

The Deployment Platform of the Future

Migrating Operational Databases to the Cloud



SingleStore

Table of Contents

Introduction	3–5
Converging Workloads and Innovations	
Benefits and Challenges of Migrating to Cloud	6–8
Cost, Flexibility, and Scalability	
Security	
Data Impedance Mismatch	
Migration Strategy	9–13
Evaluating and Choosing the Workload	
Planning the Migration Journey	
1. Defining your Scalability Requirements	
2. Ensuring App Compatibility	
3. Defining your App Availability and Disaster Recovery Strategy	
4. Securing your Data Assets in the Cloud	
Proving Your Migration Is Successful	
Commonly Used Migration Strategies	14–16
New Cloud Host (aka “Lift-and-shift”)	
New Cloud Platform (aka “Lift-tinker-and-shift”)	
Renewed Application	
Introducing Managed Service by SingleStore	17–21
Benefits of SingleStore Managed Service	
Choosing the Right Cloud Database	22–23



Introduction

If data is the fuel of the digital economy, then databases are the engine that propels you forward. To keep pace with the enormous (and growing) volumes of data and ever more demanding apps and users, databases have undergone a massive transformation.

And the next wave of transformation facing databases is exploiting the deployment platform of the future: managed services in the cloud.

The database shift to the cloud has taken the market by storm in recent years. More and more database deployments and innovations are increasingly cloud-first or cloud-only. **Gartner¹ predicts that by 2022, 75% of all databases will be deployed or migrated to a cloud platform.**

Cloud databases are growing at such an unprecedented pace for many reasons. In the cloud, there is no need for technical professionals to worry about installing, configuring, managing, upgrading, or tuning the database and its underlying infrastructure—an activity that is now outsourced to the cloud database vendor. With a cloud database, technology professionals can focus on where they add the greatest value—building applications and analytical systems with greater speed, agility, and scale for today’s digital economy.

Data has been steadily moving to the cloud for over a decade, with different industries adopting cloud at different rates, and the type of data moving aligned with the maturity of the cloud itself.

The first set of data to move to the cloud was application data; as developers moved their development and deployment tools to “as-a-service” platforms, the databases they preferred to support those applications moved with them. This type of data living in the cloud was low risk and high reward. The increased developer agility and trouble-free management far outweighed the loss of control or security fears. The next wave of data moving to the cloud in the past few years has been aggregated historical data, typically in data warehouses. The tremendous cost savings of moving to lower-cost cloud infrastructure provides ample benefit to shift what is typically medium priority data used by a relatively small number of users.

The next wave of data to move en masse will be the “crown jewels” of the data in an organization: operational data. The mix of current (often streaming) and historical data actively used for core systems and mission-critical operations is beginning to shift to the cloud, thanks to the onset of improved cloud database capabilities and the flexibility and control of multi-cloud and hybrid cloud deployment options. We are just beginning what is about to be a massive migration of our most important data and workloads.

Moving an operational database to the cloud can seem quite complex, and many organizations might not know where to start. But, it doesn’t need to be like this. Whether you’re developing a comprehensive cloud migration strategy for your enterprise, or simply moving a single application to the cloud, a well-thought-out and well-executed migration plan can help you eliminate many of the unknowns and simplify your data migration efforts.

This whitepaper is written to serve as a quick guide to decide the ‘when and what’ factors involved in moving operational data to the cloud, and how a new, modern, cloud-native, fully managed database service called SingleStore Managed Service may be the ideal choice for many organizations.

Converging Workloads and Innovations

Historically, cloud databases were thought to be suitable only for running simple application workloads, and cloud database adoption was largely associated with a few cutting-edge startups. Over the years, hardware improvements in compute, storage, and networking were able to light up newer data-tier capabilities in the cloud, with promises to deliver higher scale and performance than was previously possible.

Modern cloud databases are able to leverage the latest and greatest in hardware innovation—the newest microprocessors, arrays of GPUs, the most optimized RAM and SSD storage, petabit networking, and even programmable hardware if needed. Across industries, the enterprise is adopting cloud databases to serve their daily business needs with a variety of different workloads including transactional, data warehousing, operational analytics, as well as machine learning and artificial intelligence.

This means that modern cloud databases need to fuse operational, analytical, and machine learning capabilities into the same set of cloud resources to accelerate time to insight. To reduce latency, they need to place resources at the edge, closer to users, and leverage other core capabilities like auto-tiering, caching, and other capacity optimizations to get the best performance and reduce latency.

Apart from the hardware modernization of cloud platforms, the data volumes produced by today's digital apps are also growing at an unprecedented pace and most often require instantaneous processing. With streaming and historical data used to power everything from real-time dashboards to instant decisions for apps and experiences, there is additional pressure on legacy database systems to meet all of these new requirements.

The rapid increase in the volume of incoming data has spurred many companies to consider moving to big data solutions such as Hadoop and other cloud data warehouse products to solve for this use case. Quite a few have already made the move.

However, this approach cannot deliver on the fast analytics promise, since these solutions were not built for large-scale analytics. Many of these solutions also lack SQL support, or need data to be in proprietary formats that are hard to transform into and out of, inhibiting analytical capabilities such as ad-hoc querying and the use of standard business intelligence (BI) tools.

Is there a solution? We suggest that you rethink your approach to your data technology investments, and reduce both complexity and cost with a move to a cloud-based service that provides speed, easy and low cost scale, and traditional SQL capabilities, all in one package.



Benefits and Challenges of Migrating to Cloud

Every business has some unique workloads and concerns, and there are some common challenges that come up when considering migrating to the cloud over traditional deployments.

Cost, Flexibility, and Scalability

Without any doubt, when it comes to getting more scalable data infrastructure at a lower cost, the cloud is the clear winner. In the cloud, capacity planning is virtually eliminated, and the cloud provides scalable infrastructure that can be easily managed. The absolute cost of using the cloud is likely to amount to less than the cost of buying, supporting, maintaining, and designing the on-premises infrastructure needed for your applications to run, especially when compared to the previous generations of on-premises database software.

In the cloud, patches and upgrades are hassle-free and outsourced to the cloud provider, so you can focus on optimizing your existing IT investments. With its resilient architecture across multiple availability zones, security features baked into every service, as well as more consistent logging and monitoring tools, the cloud also provides peace of mind.

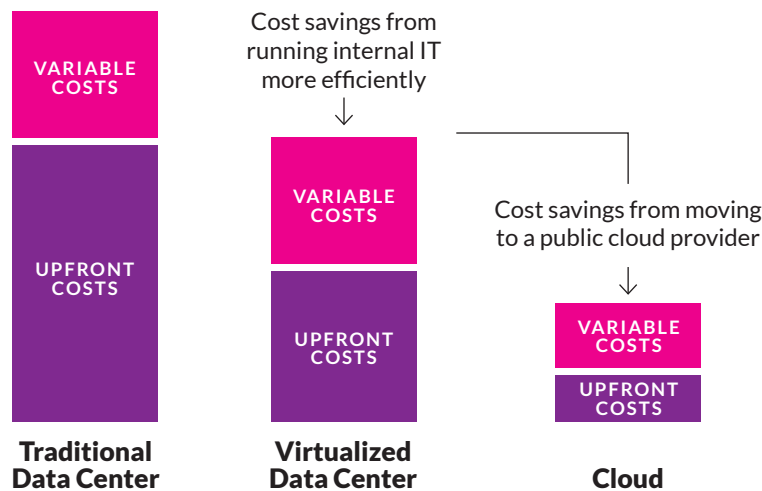


Figure 1. Cost savings from moving to the cloud

Organizations that have successfully taken advantage of the cloud are not simply enjoying incremental improvements. For example, legacy batch processes that were used for historical data analysis now have real-time execution requirements, with the use of predictive analysis and AI. Along with analytics, ML and AI, cloud-native technology is also growing at an unprecedented pace. The most successful cloud-native strategies are achieved with the use of platform-independent, software-defined infrastructure that can be efficiently managed through automation and commoditized tooling.

The right cloud database can support a variety of workloads and ensure compatibility with existing applications and tools.


Security

Simply put, data security concerns can outweigh the ease of access and flexibility that cloud databases provide. Despite the fact that major cloud providers usually offer high-grade security on both the hardware and software end, insufficient security controls or the lack of security-by-default designs can still compromise your data. Enterprises might be tempted, especially during the migration process, to lower security barriers in favor of giving speedy, painless data access to its developers from the cloud; doing so presents obvious security risks that can be challenging to discover and fix later (and perhaps too late).

Data Impedance Mismatch

Enterprises that work frequently with non-standard data types in their applications face a challenge due to data impedance mismatch issues. Some cloud database vendors might not offer support for all of the required data types “Out of The Box,” and in these cases, your applications might need revisions and/or the adoption of additional cloud databases to get them working in the cloud. This could include code rewrites, or complicated data export, transform, and load (ETL) operations. Additionally, the lack of a solid knowledge base and support from cloud vendors can make the migration effort more stressful and difficult to accomplish.

Fortunately, these issues can be overcome by developing and implementing an effective migration plan. So, what does an effective migration plan look like? The next sections lay out a framework that can be adapted to specific instances.



Cloud Migration Journey

Cloud migration is a journey, and there is no one-size-fits-all solution. As we mentioned previously, for some, migration to cloud might mean moving just a standalone application. For others, it may mean moving a broad cohort of applications and analytical systems. While there are several common elements found in each successful migration, taking a strategic, intentional approach to migrating your data tier to the cloud can alleviate many of the technology's core challenges.

Following these three steps can help your company navigate the transition successfully.

Evaluating and Choosing Workloads that Benefit Most from the Cloud

Once you've decided to migrate to the cloud, consider re-thinking why you want to embrace the cloud in the first place, and the goals you want to reach. Even though reacting quickly to marketplace pressures (and strategic decisions from management) is critical, racing ahead without a clear understanding of all the motivating factors can lead to a rushed move that may ultimately fail to yield the results you were looking for. Before you start migrating your workloads to the cloud, evaluate and assess your environment thoroughly to decide what would benefit most from being migrated, your readiness to support a migration of that asset, and what impacts might arise from the migration.

It all starts with understanding the state of your on-premises workloads, segmenting and prioritizing them, and then creating a migration plan for the workloads that make sense for your business goals, capabilities, and ability to support the change. Typically, applications and databases which already need to be modernized, enhanced, or improved to support new requirements or increased demands are a good first choice. You will already be making substantive changes to the code base and supporting infrastructure, so the impact of required changes to fit the cloud are lessened, while the benefit of instantly available and elastic services helps with development, testing, and deployment. You might also want to consider those apps with highly variable throughput. Applications that experience large seasonal changes in volume, for example, can benefit from the compute elasticity the cloud provides.

When it comes to databases, moving between databases that support the SQL standard is not generally a complex process, since most of the data types will be common, and simple mappings can be found if not. When you move beyond relational databases, the mappings are significantly more challenging, and might require a full application rewrite or extensive ETL tasks. Choose a database that will fit both the data types and structure you are already using on-premises, as well as the skills of the people who are responsible for using or managing those databases.

The process of migrating data infrastructure to the cloud should be carried out in an agile, flexible and iterative way.

Planning the Migration Journey

A common misconception about planning your cloud migration is that it has to be a one-time trip. In reality, the process of migrating data infrastructure to the cloud should be carried out in an agile, flexible and iterative way. It all begins with defining a baseline and an ongoing plan for improvement. To do this, characterize how your apps are behaving on-premises, and ensure that these characteristics will be achievable in the cloud. Here are a few things that you might want to think about as you define your journey to the cloud in steps.

1

Defining Your Scalability Requirements

The elasticity of the cloud can enable you to get more scale for your data infrastructure, and ultimately for the apps that depend on it. Even though it is hard to predict when your app will see a surge in traffic, knowing the upper scaling limits of your application will make it easier to apply scale-out operations, and minimize under- or over-provisioning of cloud resources. Cloud platforms have scalability features that you can take advantage of to scale compute and storage resources, resulting in low-cost, economical storage for large data volumes. This can be extremely useful when, for example, you know that an application always sees peak utilization in the morning, but low utilization at night.

2

Ensuring App Compatibility

Before moving your data to the cloud, it is important to ensure that the data tier in the cloud is compatible with the data types required for your app, and for your data tier to fully support the SQL query language. (These requirements can work against each other, because many of the cloud-platform-specific relational databases tend to offer excellent SQL support, but work with a narrower range of data types than NoSQL databases and other kinds of data stores.) With consolidation as one of the top reasons for moving to the cloud, your data tier should integrate both structured and unstructured data, and ideally would bring them together in one database to avoid creating new data silos. To migrate your data to the cloud, you should also consider taking backups on-premises that can be shipped to the cloud and restored, to initially populate the database. Additionally, you might also want to configure ongoing data replication between the on-premises instance and the cloud instance, offering long-term control and the flexibility to switch providers later.

3

Defining Your App Availability and Disaster Recovery Strategy

Whether you're running your apps on-premises or in the cloud, service disruptions can happen at any time. Your network could have an outage, your latest application push might introduce a critical bug, or the data center might be hit with natural disaster. Before you move to the cloud, you should define your availability requirements for your application, and have a robust, well-tested disaster recovery strategy.

First, you need to evaluate the data needs of your app to determine whether they simply need only a backup taken, or if they need to be included in the disaster recovery strategy. By replicating your data across multiple cloud providers or regions, you can ensure that your apps are always online, and that your data is always available, even if a data center goes down completely.

4

Securing Your Data Assets in the Cloud

When it comes to sensitive data workloads running in the cloud, it is important to leverage built-in security controls to meet regulatory and compliance requirements. However, knowing whether a database is securely configured, who can access it, and where does the sensitive data reside is challenging for most organizations.

When evaluating a cloud database, database security must be taken into account from the very beginning. This means that you must carefully evaluate the database technology and security posture of the database vendor to ensure that security best practices are used by default, and that you have all the necessary tools for rolling out and utilizing the latest security controls including role-based access control, industry-standard password protection policy, and strong data encryption.

Proving Your Migration Is Successful

It goes without saying, but proving that your migration is successful is needed to confirm that everything is actually working as intended. If you've moved your existing application without many changes, you should check that SLAs (Service Level Agreements) of your application are still within the specified limits before completing the migration. In the case of unpredictable traffic spikes, you need to ensure that your data infrastructure can scale up and down dynamically to cope up with the additional load. If you've replicated the app data across multiple regions, you should be able to see improved response times for your applications and a smoother user experience. All migrations are driven by a business goal, so whether the reason to migrate was cost reduction, or just to gain additional developer agility, you must compare the completed deployment against the original deployment to verify that the business objectives are actually met.



Commonly Used Migration Strategies

Now that you know what to look for in your cloud migration journey, in this section, we will cover some commonly used cloud migration patterns.

New Cloud Host (AKA “Lift-and-shift”)

For many applications, moving to a new cloud host, also known as “lift-and-shift,” is a common migration pattern. It involves moving applications from an on-premises host to a new cloud host, without the overhead costs of maintaining on-premises servers after migration is complete. Imagine “copy and paste”-like migration.

With the lift-and-shift pattern, many organizations can continue to operate with minimal service interruption, avoiding the high cost of rebuilding and redesigning their apps. This also means, however, that applications may not be optimized for their new host environments in the cloud. When applications are shifted, they don’t necessarily have to “freeze”; you can still refactor or revise them until everything is “just right”.

New Cloud Platform (AKA “Lift-tinker-and-shift”)

A related approach, also known as “lift-tinker-and-shift,” is to move to a new cloud platform, but with some changes to the app. Compared to the previous pattern, in this pattern, the core architecture of the application remains the same, but a few improvements to the app and cloud optimizations are done to get added advantage. For example, an organization might replatform the application to use a new cloud database or standardize on a particular OS in the cloud.

Renewed Application

In this pattern, the application is “renewed”: completely rewritten from scratch for the cloud platform. It is much more complicated to do, since the organization needs to ensure that all the data and business logic are migrated successfully. On top of that, retraining of staff might also be required, since the new application could function quite differently from the previous one. However, the rewritten app has much more opportunity to be re-considered and optimized so as to get the most out of the cloud environment.

Choosing a migration strategy is easy. Designing and implementing it is the challenge. Many migration projects run into delays, cost overruns, and other problems that can disrupt your business.

In most cases, it's not the application, but the data tier that is not compatible.

What if your database in the cloud could be truly identical to what you have on-premises so you can easily “lift-and-shift” or “lift-tinker-and-shift” your app to the cloud, while still getting all the benefits the cloud can offer?



Introducing Managed Service by SingleStore

SingleStore
can ingest
many millions
of events per
day, with
support
for ACID
transactions

SingleStore is an operational database built for performing both transactions and analytics to support the demands of modern applications, analytical systems, and ML/AI at scale. SingleStore uses a cloud-native, distributed architecture to deliver maximum performance and elastic scale. Note that “cloud-native” does not mean “cloud-only”; in fact, “cloud-native” infrastructure and apps are completely flexible, being able to run on any cloud or on-premises. SingleStore does so by offering both multi-cloud and hybrid options, ranging from a database-as-a-service, to Kubernetes-based hybrid and private deployments, to traditional on-premises installations on VMs or commodity hardware.

SingleStore can ingest many millions of events per day, with support for ACID transactions, while simultaneously supporting analytics, applications, machine learning and AI queries on trillions of data rows. SingleStore can support running transactional and analytical workloads under high concurrency, all while still supporting the standard ANSI SQL query language.

If you’re already familiar with running SingleStore on-premises, you can now enjoy the ultra-high performance and elastic scalability of SingleStore in the cloud. SingleStore’s database-as-a-service offering is called SingleStore Managed Service. SingleStore Managed Service gives you the full capabilities of SingleStore without the operational overhead and complexity of managing it yourself. SingleStore Managed Service provides a resilient database with cloud-agnostic deployment support on AWS and Google Cloud Platform (with support for Azure and others to come). With SingleStore Managed Service, cluster provisioning, cluster management, deployment, upgrades, alerting, and troubleshooting are all handled by SingleStore. This greatly reduces operational expenses, by shifting the database administration (DBA) tasks needed to operate your database from your organization to SingleStore. Just as importantly, SingleStore is offered at a price point dramatically lower than traditional database vendors like Oracle or SAP, and our ultra-efficient query engine means that operational costs for SingleStore tend to be lower than the proprietary offerings from the cloud service providers.

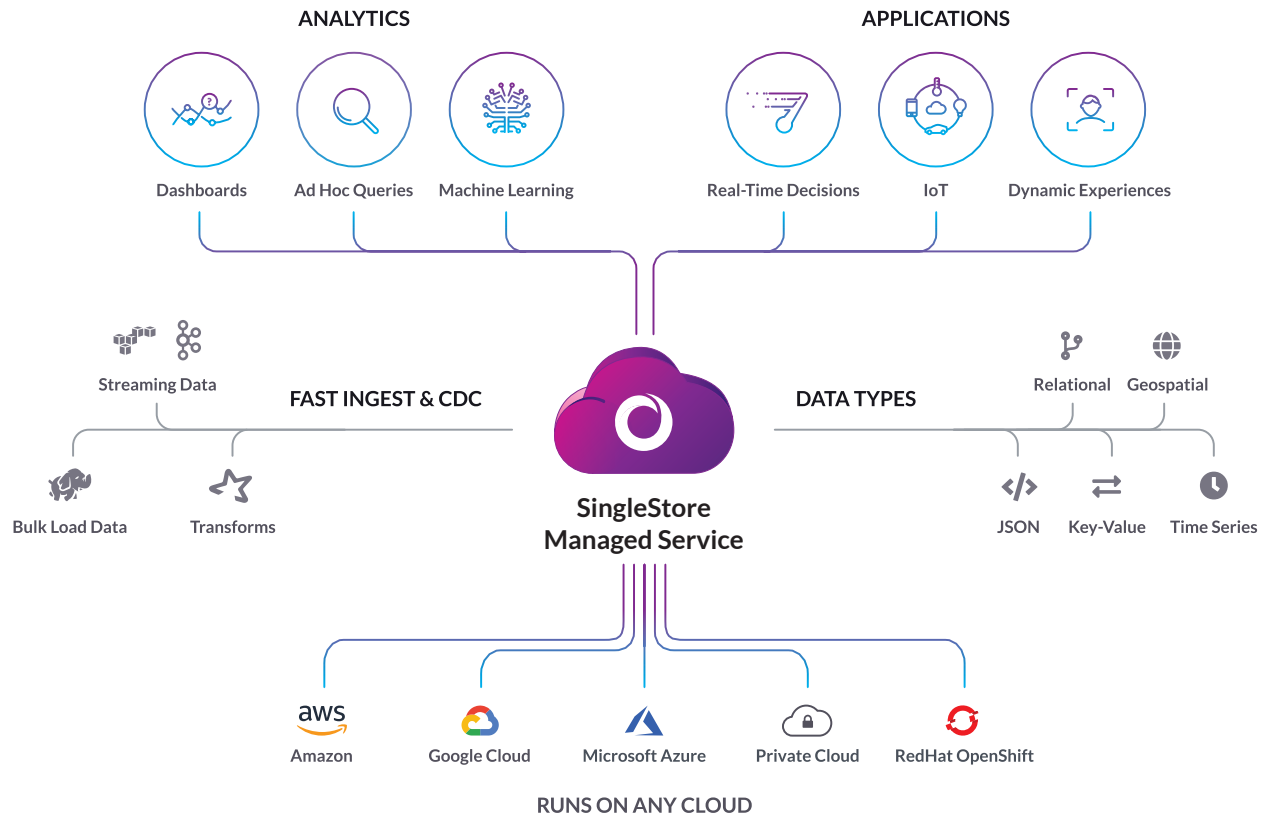



Figure 2. SingleStore Managed Service Managed Service

Benefits of SingleStore Managed Service

SingleStore Managed Service automatically backs up your data daily, with a retention period of seven days. SingleStore Managed Service runs in high availability mode, so you always have a live copy of your data, and SingleStore provides data restore services as needed. To meet regulatory compliance requirements, Managed Service supports client connections that are encrypted using TLS, and data-at-rest encryption.

Customers using SingleStore Managed Service are responsible for the logical management of their data including schema design and implementation (DDL), index and query tuning, assigning proper security permissions, requesting a database restore if needed, and requesting an increase or decrease in cluster capacity.

Depending on the needs of their application, customers have several options with SingleStore Managed Service. They can either opt to use dedicated, reserved resources if they need strong isolation guarantees, or choose to go with an on-demand model (running side-by-side with other tenants), where cluster usage is calculated hourly and billed monthly.



Effortless Deployment and Management

As we have all come to expect from cloud services, deployment and upgrades are built in. With SingleStore Managed Service you get the full benefits and capabilities of the SingleStore data platform without having to worry about deployment, management, or maintenance. There's no need to rack servers, script deployments, or manage VMs.



Avoid Cloud Lock-In Through Multi-Cloud Flexibility

Managed Service is available today on Amazon Web Services and Google Cloud Platform, and availability on Microsoft Azure is next on the development roadmap. SingleStore operates exactly the same whether deployed on-premises on bare metal, on-premises on cloud infrastructure, using the SingleStore Kubernetes Operator, or within the Managed Service service. You can use SingleStore to support a broad set of operational and analytical use cases, allowing for a simple, single platform across applications, analytical systems, and cloud deployments.



Superior Total Cost of Ownership

Compared to both legacy databases and proprietary databases from the cloud service providers, SingleStore Managed Service offers lower total cost of ownership (TCO). SingleStore offers high performance, scalability, ANSI SQL support, and the ability to replace traditional databases like Oracle Exadata and SAP Hana, at a fraction of the cost. When compared to the proprietary databases offered on Amazon Web Services and Google Cloud Platform, SingleStore's unique architecture and high-performance query engine mean that many operational analytics workloads run with far less resource consumption, offering significant cost savings. And costs for SingleStore Managed Service are very similar to costs if you use the SingleStore database in the cloud directly, and manage it yourself—but with SingleStore Managed Service, the costs and hassles of managing the infrastructure that supports the SingleStore database are included in what you pay for the managed service offering.

Choosing the Right Cloud Database



You need a database in the cloud that is fast, robust, and scalable for your operational and analytical workloads.

When it comes to migrating your workloads to the cloud, not all databases are the same. You need a database in the cloud that is fast, robust, and scalable for your operational and analytical workloads. SingleStore Managed Service is a proven, fully managed database service that differentiates itself from the rest by delivering high scale and performance for your cloud workloads, all while providing the same experience across multiple cloud providers and on-premises environments to provide greater flexibility and control.

It's critical to choose the right cloud database when moving your applications to the cloud—from saving your organization money and resources, to reducing complexity and enabling more time for innovation. You can start to move your applications to the cloud with confidence using SingleStore Managed Service. As more and more workloads shift to the cloud, the time to migrate your apps to the cloud is now, and we hope you give SingleStore Managed Service a try.

SingleStore Managed Service has quickly made it to be the heart of many enterprises, enabling smarter applications and new business insights, through which your company can gain competitive advantage at a lower cost. With SingleStore Managed Service now publicly available, we hope that you can experience it for yourself, and share your success story.

**Test-drive SingleStore Managed Service at
singlestore.com/managed-service**



©2019 SingleStore, Inc.