

Serverless NoSQL Database

Efficiencies and Cost Savings from DataStax Astra on Google Cloud Platform

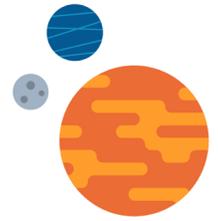
Today's data-driven enterprises are launching new business-critical applications at an incredible pace.

Increasingly, these applications must meet the market's demand for a smart "next best action" at the point when a user would find it most helpful. The only way enterprises can leverage data as a strategic asset, however, is to integrate data reliably and efficiently into business applications across the entire organization.

From an engineering and operations perspective, building and deploying the data infrastructure necessary for modern software is becoming a challenge. That's why a streamlined distributed database solution deployable to serverless environments is the way forward for many organizations.

In this whitepaper, we'll discuss the need for a modern NoSQL database solution like Apache Cassandra™, and how DataStax Astra can streamline database deployment and management on Google Cloud Platform. We'll also cover why serverless Cassandra is a game changer because of the many benefits it provides to organizations.

Enterprises Need a Modern Database Solution



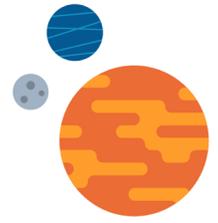
The overwhelming majority of today's enterprises are shifting to more data-driven processes, yet the relational databases that these business applications rely on can't achieve the scale and performance necessary.

There is also an increasing need for enterprises to build cloud-native solutions which enable the ability to focus on business outcomes over managing infrastructure. True cloud economics leverage the cloud to lower costs, has become essential for enterprises who are stepping up their efforts to modernize.

Legacy databases were suitable for many relational use cases over the past decades; however, such databases struggle to provide modern capabilities that enable accelerated development and faster go-to-market. Today's business-critical applications require a data layer that supports multiple data models, zero downtime, infinite scale, developer APIs, and integration with cloud services that enable analytics, event streams, and more which is a challenge for legacy databases. Architected with a primary or "master" database, legacy databases are difficult to scale and have performance bottlenecks.

At the broader level, relational databases were also not designed to handle data that pours in from mobile devices, sensor networks, retail systems, and telecommunications call-routing systems. This fast data is becoming a major challenge for data-driven organizations. Handling this fast data requires a database capable of handling massive, distributed data at speed, high-scale streaming technologies that can deliver events as rapidly as they occur, and analytics at the point of interaction.

Why Apache Cassandra?



Apache Cassandra is an open-source, NoSQL, database management system that's designed for handling large volumes of data with low latency and zero downtime.

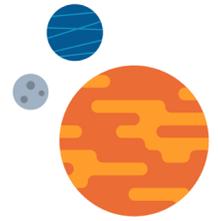
For example, Cassandra's masterless architecture allows the system to distribute read and write requests across all the nodes within a cluster. Since every node is a peer, any node can facilitate read and write requests. This enables responses in milliseconds even with large volumes of data.

Moreover, Cassandra is a NoSQL database designed for decentralized deployments, which enables the system to scale horizontally more easily than relational databases. Linear scalability means that for every node added to a Cassandra cluster, there's an equal increase in throughput or performance. Its ability to scale out quickly and handle concurrent read and write transactions allows Cassandra to outperform most relational database solutions.

Cassandra's masterless architecture also ensures that there's no single point of failure. If the "master" database in a relational system goes down, there would be a loss of write transactions during failover while a new "master" database is promoted. Cassandra replicates data consistently across multiple nodes using a hashing algorithm, so the system can continue operating without any interruption by rerouting the write traffic to any other node in the cluster.

These capabilities result in a modern database solution with high performance, 100% uptime, and consistent availability of data. Open-source Apache Cassandra is powerful; however, there are challenges that enterprises face to using this technology for modern data applications.

The Challenges with Apache Cassandra



While there are distinct advantages for adopting Apache Cassandra, the organization must overcome the administrative burden, development complexity, and increased cost. Let's take a closer look at each of these potential challenges.

➔ Administrative Burden

Cassandra's unique masterless architecture provides horizontal scalability that enables enterprises to manage millions of transactions. However, it does come with administrative challenges when it comes to maintaining, tuning, and monitoring large clusters with hundreds of nodes.

In addition, the overwhelming majority of database administrators (DBAs) and site reliability engineers (SREs) are more familiar with relational databases than Apache Cassandra or other NoSQL database solutions. That means it's challenging to find the talent necessary to deploy and maintain Cassandra, and implementations by self-taught internal staff could lead to consistent and fragmented use of the technology. The promise of 100% uptime depends entirely on having the knowledge to properly size, configure, patch, and upgrade Cassandra.

➔ Development Complexity

In addition, most organizations lack the expertise required to develop applications for Cassandra. Working with the database system requires knowledge of Cassandra Querying Language (CQL), data modeling, and installing drivers. This can slow development velocity, which is particularly detrimental as developers continue shifting to a more agile and iterative software delivery approach.

Another major challenge is that Cassandra querying requires the use of the primary key and doesn't natively support secondary indexes. That means developers cannot create queries to SELECT data using a field other than the primary key as they might with relational databases. As a result, Cassandra works best with denormalized data which means duplicating data to improve performance. The unique needs of Cassandra often make development more complicated than more familiar database technologies.

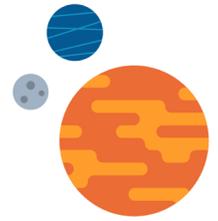
→ Cost

Most organizations that choose Cassandra see high availability as a primary goal. That means they often overprovision clusters to ensure a certain service level is achieved. Overprovisioning is common with new applications because the traffic and application requirements are not well known.

This complexity also makes budgeting more difficult, especially for event-based workloads that remain idle for much of the time. Event-based workloads include payroll processing or compiling monthly financial reports, which may be resource-intensive for short periods of time. Most organizations will err on the side of caution and ensure there's a sufficient buffer of computing resources for sudden changes in traffic or usage. At the same time, operations teams can also be slow or hesitant to scale back underutilized computing resources because scaling down can potentially result in downtime. In both of these scenarios, organizations are paying for unused cloud resources.



Introducing DataStax Astra



DataStax Astra is a cloud-native Cassandra-as-a-Service solution that simplifies the deployment and ongoing maintenance of Apache Cassandra.

Deploying and operating Cassandra requires considerable domain knowledge, but Astra makes it possible to rapidly adopt the database solution and efficiently manage workloads. Astra can decrease the deployment time of Apache Cassandra from weeks to minutes and is available through the [Google Cloud Platform](#) Marketplace or directly from DataStax website.

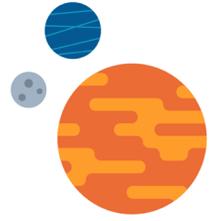
Apache Cassandra can be seamlessly deployed on [Google Cloud](#), with a cluster installed in a single region and across multiple zones. This ensures Cassandra nodes are distributed amongst Google's numerous data centers for maximum data redundancy. Astra eliminates the administrative overhead with automated infrastructure backups, patches, and upgrades. Customers with Google Cloud spend commitments can now use those funds to purchase DataStax via the Google Cloud marketplace.

The [Stargate](#) open-source data API gateway is integrated into Astra. It eases the burden for developers because they can use modern APIs, such as REST, GraphQL or schemaless JSON documents to integrate Cassandra with data-driven applications. These are technologies developers work with on a daily basis, so they won't need to learn CQL or install client drivers. The Document (JSON) API even enables developers to create applications without any upfront data modeling. Instead of learning a new querying language or building new endpoints, development teams can begin working with Astra databases immediately.

In addition, DataStax has created a storage-attached index (SAI) for querying Cassandra databases using a secondary key. SAI allows developers to build applications more efficiently because they can use secondary keys to query the Cassandra database, similar to how they'd work with a relational database.

Finally, Astra is a database as a service (DBaaS) solution, meaning there's no capital expense burden or upfront costs. Enterprises can simply pay for the resources they want, and leave database administration details like deployment, patching, backups, compaction, and scaling, to DataStax's team of experts.

DataStax **Astra** on Google Cloud - Better Together



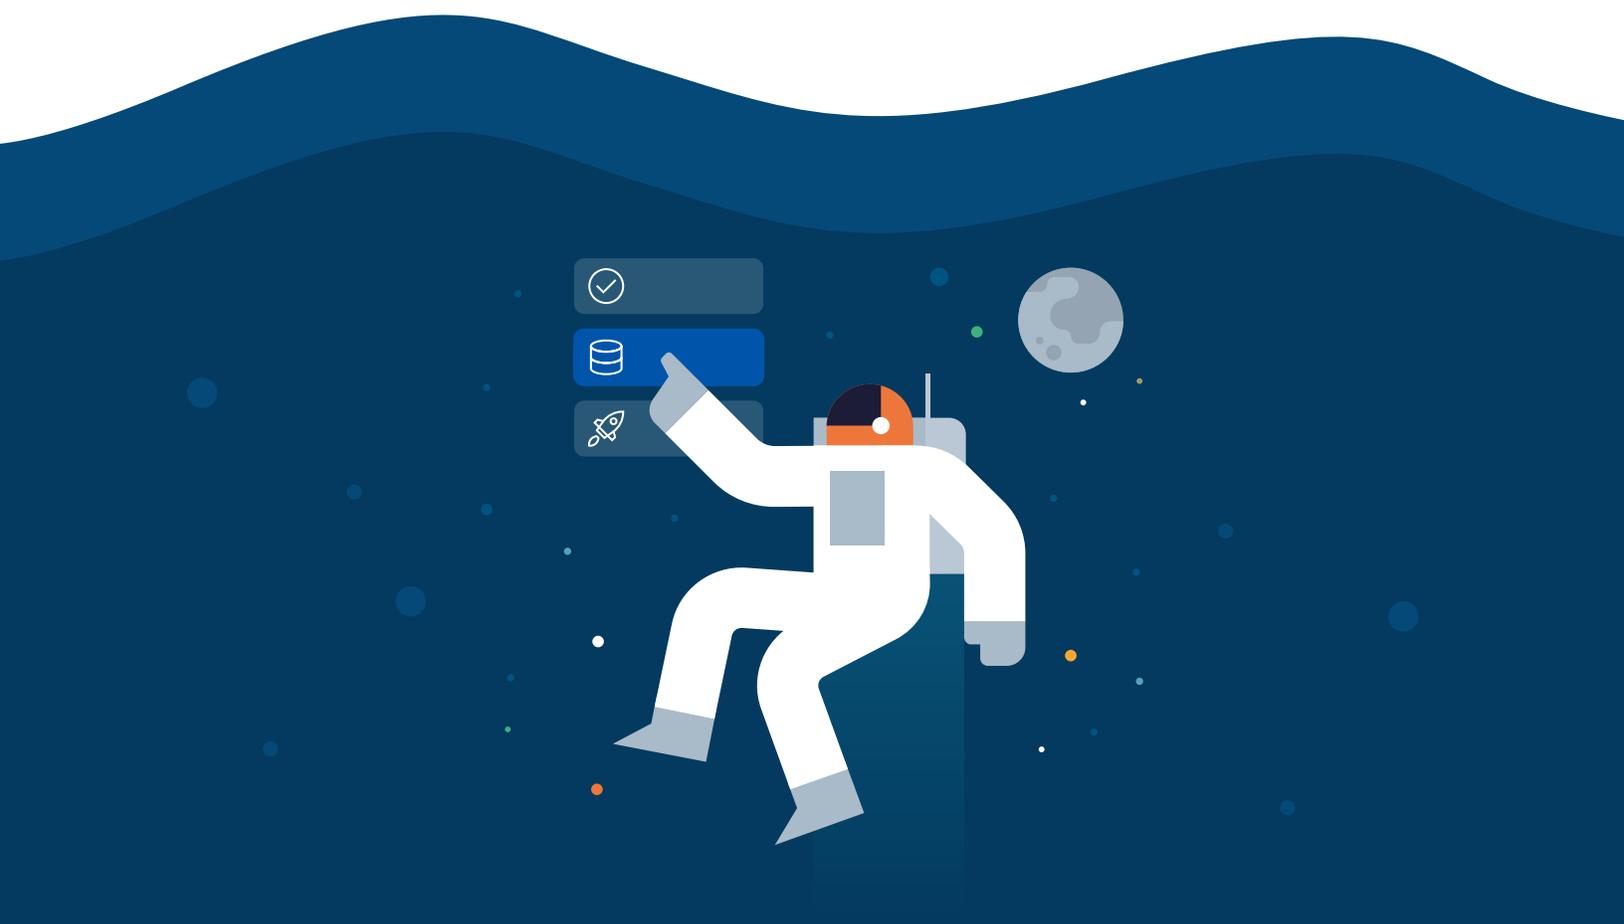
While many cloud providers are competing in the market around capabilities, regions, and a multitude of services, Google Cloud is paving its own path into the heart and soul of enterprises.

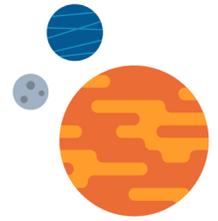
Google Cloud offers great services, but it also offers them in a manner that seems inclusive because of the unified network and service ecosystem. This unique ability of Google Cloud creates a unique experience for enterprises who are looking to modernize with Astra to get true cloud economics with serverless and dynamic scaling.

Here are few critical advantages to enterprises:

- ➔ **Network Excellence:** Enterprises now have the ability to develop and deploy their applications with speed using APIs available in Astra on Google Cloud that leverages Google's unified network that is powered by faster cable network that provides speeds up to 10Tbs.
- ➔ **Multi-region reliability:** The capability to have the same underlying network for multi-region in Google Cloud provides exceptional user experience for end-users of Astra.
- ➔ **Migration Benefits:** Astra users are migrating from legacy platforms and migration on Google Cloud is simplified to ensure success for developers and operators. This allows quicker migration and ensures faster GTM.
- ➔ **Unique AI/ML/Big Data Services Blend:** Astra handles the transactional fast layer which is capable of integrating with Google services which give enterprises using Google Cloud an unfair advantage when it's tied to Astra.
- ➔ **More Bank for the Money:** Astra provides the unique ability to dynamically scale up or down depending on the application workloads that are deployed on Google Cloud. Google Cloud and Astra together enable enterprises to handle economics of scale more efficiently with services in Google Cloud.

- ➔ **Security Superiority:** Google's process-based and physical security investments for mission-critical applications that would run on Astra provide superior security and data protection that is essential to Enterprises.
- ➔ **Continuous Product improvements with Zero Disruption:** Both Google Cloud and Astra teams deliver manageable improvements in a continuous stream over batch improvements leading to operators seeing zero disruption in existing enterprise applications and services.
- ➔ **Developer Freedom Centricity:** Google Cloud and Astra teams are focused on providing developers the tools they are familiar with such as REST endpoints, APIs etc to accelerate development. More and more services and API integration points are available to developers on Astra making it easy to build applications and integrate with services in.



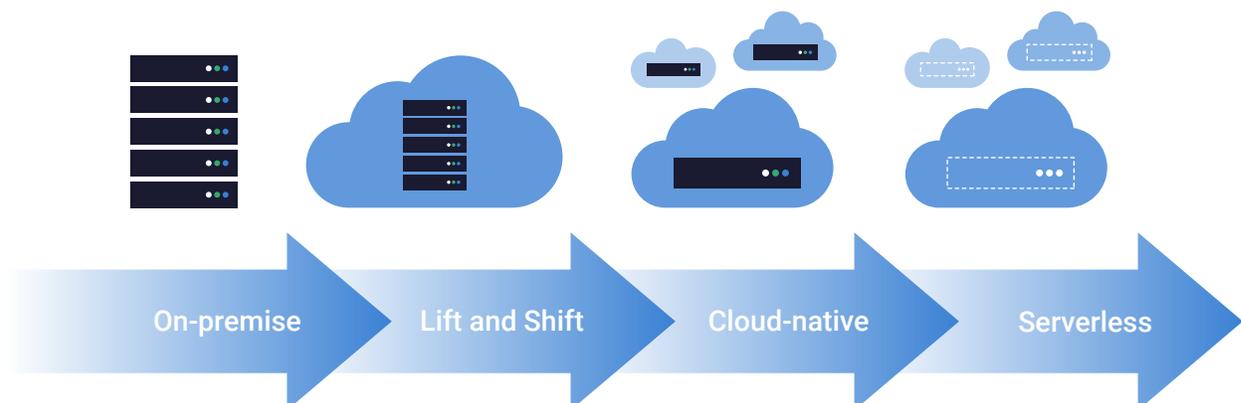


Serverless computing is a cloud model where developers can build and deploy software without worrying about the underlying infrastructure. “Serverless” doesn’t mean there isn’t a server, but that the server is instantiated on-demand at run time and destroyed or scaled to zero when not needed. Google Cloud and other serverless providers handle capacity planning, configuration, maintenance, scaling, and more for customers. Most enterprises are familiar with serverless compute; however, many organizations are choosing to deploy their data infrastructure using a serverless architecture as well.

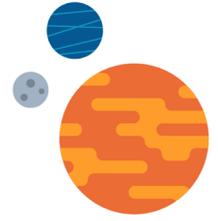
More specifically, the convenience and cost-efficiency of serverless have led to adoption across many different data infrastructure components from ETL to databases and data warehouses. These types of event-driven use cases – where processes are idle most of the time and require bursts of computing and storage when they’re triggered by an action or event – are ideal for a serverless pricing model. That said, there are many application stacks that run all of their underlying components using serverless, from databases and backend services to front-end user interfaces.

For these reasons, serverless is a natural step forward as enterprises continue to modernize their infrastructure. The shift from on-premise to cloud has accelerated rapidly, and many applications are now cloud-native, which means they’re developed in the cloud and deployed in the cloud. The convenience of on-demand computing resources could drive organizations to shift their entire stack to serverless in the near future.

Evolution of Cloud Architecture



Apache Cassandra Goes Serverless



DataStax has recently launched its own serverless solution built atop Cassandra, so organizations can unlock the power of infinitely scalable and on-demand computing resources for their database deployments.

Here's a breakdown of the key benefits of running serverless Cassandra, from reduced infrastructure costs and simplified capacity planning to streamlined application development.

➔ Reduces Infrastructure Costs

Serverless can reduce the total cost of ownership by 50% or more in terms of both upfront and ongoing costs. Third party validation by GigaOm concluded that serverless Cassandra saved 76% on TCO when compared with self-managed Cassandra clusters. And when choosing amongst different cloud providers, Google Cloud is the best option as its storage and throughput costs are lower than competitors. In particular, when we consider traditional Apache Cassandra use cases, which are typically write heavy, the savings with Google Cloud can be more than 25% over other cloud providers.

The elasticity of serverless greatly reduces the cost for workloads with unpredictable or infrequent capacity requirements because enterprises only pay for computing resources when these processes are running. As a result, enterprises can utilize a pay-as-you-go pricing model, by paying for the number of reads (in millions), writes (in millions), and storage and data transfers (in Gigabytes). If the database is not in use they pay nothing. If an enterprise wants to reserve database capacity they can select an alternate model that commits to a monthly spend in exchange for better pricing.

→ **Simplifies Capacity Planning**

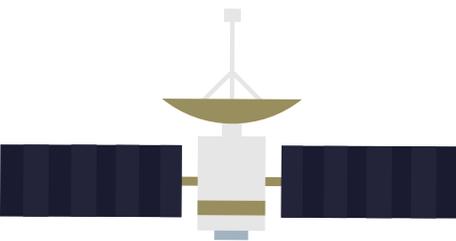
Capacity planning is always a challenge for operations teams because nearly every app has some seasonality or variable usage throughout the year, and it's also hard to predict usage for new applications. Along with seasonality and unpredictable growth, there are event-driven use cases like machine learning or data processes that require temporary capacity increases.

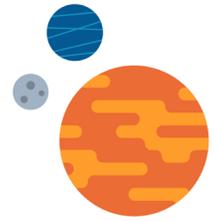
Serverless is highly elastic, so enterprises can closely align resource capacity with demand automatically. That means operations teams won't need to estimate capacity or overprovision resources for unexpected usage spikes. Cassandra by its very nature easily scales horizontally (adding additional nodes) but is difficult to scale down. DataStax has separated Cassandra into multiple microservices, the biggest being storage and compute, each of which can be scaled independently. This is an engineering feat, most databases are serverless because compute and storage are tightly coupled. As a result, DataStax Astra is the only serverless Cassandra-as-a-Service in the market today.

→ **Streamlines Development**

Since serverless abstracts away the underlying infrastructure, engineers can focus on developing innovative, data-driven software rather than worrying about databases and cloud resources. That means there's no longer a barrier between developers and cloud operations teams, and organizations can adopt a more agile DevOps approach to software development.

Astra Serverless also makes it convenient for organizations to provide developers with their own Cassandra database to use as a development environment. By only paying for reads, writes, and storage on a monthly basis, organizations can ensure development teams have sufficient testing environments without wasting resources. Astra's role-based permissions on a granular level also allow organizations to give developers access to certain clusters, databases, and tables without introducing privacy and security risks.





Modern enterprises shouldn't compromise on their data infrastructure. Instead of self-managing a disparate data stack of open-source or proprietary databases and computing infrastructure, organizations should consider a fully managed DBaaS or serverless solution like DataStax Astra.

DataStax is a leading provider of database products and cloud services, and a large contributor to the open-source Apache Cassandra project. Astra on Google Cloud ensures enterprises have a production-ready Cassandra from day one to minimize the costs and risks associated with operationalizing data. For organizations seeking to create cloud-native applications, Astra runs on Kubernetes and uses a microservices architecture and API first approach.

DataStax Astra combines the unique qualities of a NoSQL, distributed database and serverless to achieve 100% uptime, high-performance, and a reduced total cost of ownership. In fact, it's currently the only Cassandra-based serverless data solution on the market. Ensure the data for business-critical applications is always available when needed without worrying about data and cloud operations.

Serverless is the next step in the evolution of cloud infrastructure and the new default option for DataStax production and free deployment. With DataStax Astra, we're confident organizations can streamline the transition to Apache Cassandra and immediately see greater operational efficiency and cost savings.

Sign up for a free account today on [Google Cloud](#) or the [DataStax Astra](#) webpage. Your free account entitles you to 30 million reads, 4.5 million writes and 40GB free monthly.

About DataStax

DataStax is the open, multi-cloud stack for modern data apps. DataStax gives enterprises the freedom of choice, simplicity, and true cloud economics to deploy massive data, delivered via APIs, powering rich interactions on multi-cloud, open-source, and Kubernetes. DataStax is built on proven Apache Cassandra™, Apache Pulsar™ streaming, and the Stargate open-source API platform.

DataStax Astra is the new stack for modern data apps as-a-service, built on the scale-out, cloud-native, open-source K8ssandra. DataStax powers modern data apps for 500 of the world's most demanding enterprises including The Home Depot, T-Mobile, Intuit, and half of the Fortune 100.

© 2021 DataStax, All Rights Reserved. DataStax is a registered trademark of DataStax, Inc. and its subsidiaries in the United States and/or other countries.

Apache, Apache Cassandra, Cassandra, Apache Pulsar, Pulsar, and Apache Kafka are either registered trademarks or trademarks of the Apache Software Foundation or its subsidiaries in Canada, the United States, and/or other countries.