

Kubernetes for Enterprises


Real-Time Everywhere With DataStax and Google Anthos

Enterprises in industries like eCommerce and the Internet of Things (IoT) continue to move towards microservices-based architectures and containerization.

While these application architectures help organizations achieve lower latency and improved reliability, Google Anthos is the next step forward for modern distributed infrastructures.

In this whitepaper, we'll discuss Google's hybrid and multi-cloud Anthos platform for modernizing enterprise applications, why distributed databases are crucial for this low-latency, active everywhere architecture, and how DataStax Enterprise fits into the mix.

What is Google Anthos?



As enterprises move towards containerized deployments, orchestration across decentralized infrastructure has become a challenge.

Kubernetes—the open-source orchestration platform of choice for many organizations—allows operations teams to declaratively use data center resources. The problem is that expanding Kubernetes beyond a single data center generally requires a bespoke solution.


Google Anthos aims to solve this issue by extending Google Cloud services to support multi or hybrid cloud deployments without individually managing multiple distinct Kubernetes clusters or VM deployments. That means modern applications can become truly distributed across different regions or environments to achieve greater resilience, availability, and lower latencies.

Anthos uses Google Kubernetes Engine (GKE) as the local control plane for containerized deployments while the Anthos platform acts as a central control plane for the overall decentralized architecture. Anthos and GKE together bring consistency to containerized applications deployed across various cloud environments from Google Cloud and Amazon Web Services (AWS). By deploying GKE on-prem, Google will be able to manage the upgrades and patches for Kubernetes running atop on-premise infrastructure as well.

In addition, Anthos Config Management simplifies enterprise security and compliance by allowing operations teams to roll-out policies across all Kubernetes clusters from a centralized interface. This eliminates the need for bespoke configurations so that administrators can enforce a single set of security and compliance guardrails across an ever-growing sprawl of clusters.

The Anthos platform also streamlines the development and deployment of cutting-edge applications across modern and legacy enterprise infrastructure by allowing developers to work with a single set of container images. This eliminates the need for tedious tasks to make images compatible with different Docker and Kubernetes versions.

Why Enterprises Need Cassandra For Anthos



When it comes to highly distributed applications, Anthos can streamline the management and deployment of Kubernetes clusters and virtual machines. The challenge is that most modern applications aren't stateless, which means the data layer is often a limiting factor when adopting a truly decentralized architecture like Anthos.

Most traditional databases cannot fully support the Anthos architecture because they're centralized. Relational databases usually have a primary database or node for write requests that is replicated to the rest of the secondary databases available for read requests. While these database systems were suitable during the era of monolithic applications, they're not aligned with today's microservices-based architectures. The primary-secondary approach creates a bottleneck that limits the horizontal scalability and resiliency of the data layer.

The [*DataStax Enterprise \(DSE\)*](#) distribution achieves greater performance by replicating data immediately, so that even remote clusters are in sync at the speed of the network. Other database systems use sync cycles or require complicated synchronization with all writers across the entire distributed system. Both of these approaches add time and fragility to the data layer.

Even many distributed databases aren't ideal for Anthos based applications because they operate with a primary-secondary data replication model as well. While these distributed databases can improve the availability of data through easier horizontal scaling, they also rely on primary nodes for write requests. This approach increases fragility by not allowing for regional network partitions.

The bottom line is: the Anthos architecture works best with a truly distributed and decentralized database system where every node is equal. If the data layer cannot read and write from any node, there will inevitably be bottlenecks that can impact application performance.

Consider Apache Cassandra


Apache Cassandra is a distributed database that can read and write to any node in a cluster without relying on any primary nodes. This decentralized architecture is ideal for Anthos because changes to data at individual Kubernetes clusters reflect immediately across the entire managed infrastructure. Availability is guaranteed with Cassandra's active everywhere architecture.

Along with its unique data replication properties, Cassandra is open-source and can run nearly anywhere. Most proprietary database technologies are limited to certain cloud environments, but Cassandra can be deployed to bare metal, on-premise environments, Google Cloud, AWS, and more. This freedom of choice aligns closely with the Anthos mission for managing Kubernetes anywhere.

The focus on availability and flexibility is invaluable, but Cassandra alone isn't enough. Organizations adopting Anthos should consider DataStax Enterprise (DSE), which brings enterprise-grade operational reliability to the open-source database technology.



The Benefits of DataStax Enterprise on Anthos



DataStax Enterprise (DSE) is a proprietary distribution of Cassandra that can be deployed anywhere from multiple clouds and Kubernetes clusters to bare metal and VMware infrastructure. DSE's built-in Kubernetes operator enables cloud-native Cassandra workloads that outperform other versions of Cassandra, with 50% lower latency and 4x faster scaling thanks to zero-copy streaming.

DSE is also ideal for the Anthos data layer because the database technology goes beyond Cassandra's everywhere architecture to provide additional security, monitoring, and other advanced features.

Security

With DSE, security is a high priority. That's why DSE supports encryption, single sign-on, and role-based permissions out-of-the-box. Organizations can also detect and prevent potential breaches using auditing and log scanning tools. These features ensure enterprises keep their application data secure and protect the privacy of their users.

Monitoring & Support


DSE offers the monitoring and support that enterprises have come to expect. The DSE OpsCenter provides real-time and historical performance metrics for clusters, nodes, tables, and more. In addition, the DSE distribution has been hardened by numerous Fortune 100 companies and the largest Internet applications to ensure a high level of performance and reliability.




Advanced Features

Along with security, monitoring, and support, here are some of the advanced features available for enterprise organizations:

- **Multiple Data Models:** Support for different NoSQL data models, such as key-value, wide-column, document, and graph. This flexibility allows developers to design data models that align closely with the requirements of their applications.
- **Enterprise Search:** Indexing data for enterprise search engines like [Apache Solr](#) enables full-text search, faceted search, geospatial search, and more. This can streamline the implementation of enterprise search for websites or eCommerce stores.
- **Data Analytics:** An integration with the [Apache Spark](#) analytics engine allows organizations to process large data sets. This enables enterprises to leverage Big Data analytics and machine learning to make the most of their data.
- **Real-time Streaming:** Leverage [Apache Kafka](#) for real-time streaming and distributed data pipelines at scale. This makes it easier to ingest large amounts of data from IoT sensors from any node.
- **Stargate APIs:** Reduce the learning curve for developers with powerful and flexible APIs that abstract away Cassandra-specific concepts. Developers can work with modern REST APIs, GraphQL, and schemaless JSON to quickly integrate DSE with their applications.



Anthos & DataStax in Action



While Anthos and DataStax enable any organization to adopt this new distributed infrastructure paradigm, two areas that could benefit the most are eCommerce and Internet of Things.

eCommerce

As eCommerce companies continue to become more globally competitive, low latency is a growing priority. Online shoppers won't tolerate slow websites and poor shopping experience, but they may be accessing online stores from mobile or remote networks with limited bandwidth.

Running an eCommerce website atop Anthos and DataStax can ensure low latency throughout the entire shopping process. Kubernetes clusters deployed in key regions around the world can boost page load speeds for global customers, while DSE can provide high-performance search functionality.

Internet of Things

When it comes to the Internet of Things (IoT), sensors are continuously collecting an enormous amount of data that needs to be written to the data layer at a high rate. The challenge is efficiently replicating this data layer across a globally distributed infrastructure.


Deploying Kubernetes clusters managed by Anthos closer to IoT sensors can reduce latency and lower the processing costs when collecting data. In addition, DSE can ensure this data is immediately replicated across the entire data layer so that it can be analyzed in real-time.



Want a fully-managed database solution?

DataStax Astra can also be deployed on Google Anthos through a service broker. This enables enterprises to seamlessly orchestrate containerized applications across multiple clouds using Cassandra as a highly scalable data layer. Let DataStax manage the data infrastructure so that your organizations can focus on innovation.

Choose Real-Time Everywhere



Take full advantage of the operational consistency of Google Anthos and the data replication capabilities of DataStax Enterprise to streamline modern application development and deployment. With this innovative architecture, managing distributed Kubernetes clusters across multiple environments with a single pane of glass is now a reality.

Are you struggling to maintain a fragmented Kubernetes deployment? Reach out to learn if DataStax on Anthos is the best way forward for your hybrid or multi-cloud enterprise infrastructure.

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.