

Google Cloud



The new way to build with Google Cloud databases

70+ customer stories.
Every industry. Endless possibilities.





Foreword

Thank you for interest in this eBook.

As a Google Cloud partner, we at Brilyant IT Solutions are committed to providing you with the resources and insights you need to make informed decisions about your cloud journey.

We believe that this eBook will provide valuable insights into the benefits and capabilities of Google Cloud. Brilyant is equipped to support you with seamless database modernization on Google Cloud - helping you scale operations, boost performance, and reduce complexity.

If you have any questions or would like to discuss your specific requirements, please do not hesitate to [contact us](#).

Sincerely,

Brilyant IT Solutions



Google Cloud
Partner

Table of contents

 Retail and consumer packaged goods

 Financial services

 Technology

 Startups and digital natives

 Healthcare and life sciences

 Games

 Manufacturing and supply chain

 Security

 Media and entertainment

 Telecom

 Get started with Google Cloud databases

Retail and consumer packaged goods

From storefront to fulfillment, make every decision data-driven and customer-focused.



**Industry**

Retail and CPG

Country

United States

Google Cloud databases[Spanner](#), [Bigtable](#)

Macy's manages billions of data points at scale with Spanner and Bigtable

Challenge

With 700+ stores and millions of products, Macy's generates billions of data points requiring multiple terabytes of storage. Managing this much data requires a centralized system to provide internal clients with consistent, real-time access to critical information.

Solution

Macy's migrated from on-premises operations to Google Cloud's managed database services, using Spanner for inventory management and Bigtable for its pricing system. Together, these services enhance operational efficiency, expand Macy's reach, and support seamless, real-time pricing and inventory updates across its digital and in-store channels.

Outcomes

- Bigtable delivers sub-10ms latency at p99 **for real-time pricing data** across billions of records
- Spanner **supports 10,000+ queries per second** and maintains an SLA of 99.99% uptime
- Google Cloud's managed services **enhanced the customer experience** by providing high availability and scalability for Macy's growing product catalog

[Read the full story](#)

“Our ability to enhance the performance of our operations and deliver a better experience for our customers is a direct reflection of Google Cloud managed services. The success of our partnership with Google reflects a mutual commitment to embracing innovation and imagination.”

MOHAMED NAZEEMUDEEN

Director of Software Engineering, Macy's

**Industry**

Retail and CPG

Country

United States

Google Cloud databases[Spanner](#), [Cloud SQL for PostgreSQL](#)**Additional Google****Cloud products**

BigQuery, Google Kubernetes Engine, Dataflow, Pub/Sub

Wayfair transforms its databases for speed and scale with Spanner

Challenge

Wayfair, which operates a massive ecommerce platform with more than 22 million products, needed to migrate 10,000+ Microsoft SQL Server databases from on-premises data centers to Google Cloud – without disrupting operations for their team of over 3,000 engineers, 16,000 supplier partners, and tens of millions of customers.

Solution

Wayfair began its database migration to Google Cloud with Cloud SQL for SQL Server, enabling a rapid transition of existing workloads. The company then adopted Spanner for high-throughput distributed workloads and Cloud SQL for PostgreSQL for relational applications. Wayfair also built an internal Database-as-a-Service (DBaaS) platform, making it easy for developers to provision databases on demand.

Outcomes

- **Net promoter scores from engineers have increased** by 29% for support and 41% for tooling
- DBaaS automation **reduced infrastructure management overhead**, allowing teams to focus on innovation
- Developers can **request and deploy databases in hours instead of days**, improving engineering velocity

[Read the full story](#)

“Working with Google Cloud as a cloud provider reduces our time to market to support new use cases, reduces our operational overhead, increases developer velocity, and enables us to scale at the speed of our business.”

PHIL PORTNOY

Associate Director, Wayfair

**Industry**

Retail and CPG

Country

United States

Google Cloud database[Spanner](#)**Additional Google****Cloud products**

Dataflow

Kroger enhances productivity and the customer experience with Spanner

Challenge

Kroger, one of the largest grocery retailers in the US, wanted to enhance store operations by giving associates real-time tools to optimize stocking, staffing, and store management. The goal was to improve productivity and the customer experience for Kroger's 2,800 stores and nearly 11 million customers.

Solution

Kroger worked with Google Cloud and Deloitte to develop two applications: (1) a task management app to help night crew managers prioritize stocking in real time and (2) a store management app to digitize audits for consistency. Built on an event-driven architecture, the two-pronged solution uses Spanner for real-time data tracking, Dataflow for processing sales forecasts and staffing levels, and Google Cloud's AI/ML to optimize task lists dynamically.

Outcomes

- Automated task prioritization has **improved associate efficiency** across all stores
- Digital store audits have increased consistency and **enhanced the customer experience**
- Real-time data insights help **managers make better operational decisions**

[Read the full story](#)

“Google Cloud and Deloitte brought us a technology architecture and application framework that we could implement in record time. We’re already seeing results across our stores, with associate tasks being optimized and overall productivity increasing.”

JIM CLENDENEN

VP, Enterprise Retail Systems, Kroger

**Industry**

Retail and CPG

Country

United States

Google Cloud database[Memorystore for Memcached](#)

Instacart cuts Memcached costs by 23% with Memorystore

Challenge

Instacart, a leading grocery technology platform, previously used self-managed Memcached on Compute Engine to provision, patch, and scale its infrastructure for millions of products delivered across 14,000 cities. Maintaining separate cache instances for each retailer increased operational complexity, required significant engineering time, and limited performance optimizations.

Solution

Instacart migrated to Google Cloud's fully managed Memorystore for Memcached to eliminate infrastructure maintenance. The transition, completed in two weeks with zero downtime, was supported by Terraform and Memorystore's Auto-Discovery feature, which simplified client configuration. Memorystore's optimized caching architecture also enhanced performance and reliability.

Outcomes

- **Reduced its enterprise Memcached cloud bill by 23%** using right-size instances and increasing throughput
- **Improved latency by 18.5%** on average for a subset of its ecommerce API endpoints
- Managed services **eliminated 80–100 hours of engineering maintenance annually**

[Read the full story](#)

“Had we known the full scope of benefits from switching to Memorystore earlier, we could have saved more engineering time for delivering value to other parts of our ecommerce platform.”

DENNIS TURKO

Staff Software Engineer, Instacart

Flipkart**Industry**

Retail and CPG

Country

India

Google Cloud database[Bigtable](#)**Additional Google****Cloud products**

Dataflow, Dataproc, Pub/Sub

Flipkart scales streaming analytics to handle demand spikes using Bigtable

Challenge

Flipkart, a major Indian ecommerce platform, needed a cloud-native database to replace its on-premises Apache HBase setup for its real-time streaming analytics platform – especially during its flagship Big Billion Day sales event, when transactions could spike up to 8x.

Solution

Flipkart migrated its streaming analytics platform from Apache HBase to Bigtable, Google Cloud's fully managed NoSQL database. The transition required minimal code changes, and Bigtable's built-in tooling simplified migration. By moving to a cloud-native solution, Flipkart reduced operational overhead, improved real-time analytics, and freed developers from infrastructure management.

Outcomes

- **4x auto-scaling during peak traffic** provided a smooth shopping experience
- Reduced maintenance overhead allowed engineers to **focus on innovation instead of database operations**
- **Troubleshooting time improved** with built-in metrics from Key Visualizer in Bigtable

[Read the full story](#)

“The transition to Bigtable has brought unprecedented efficiency and flexibility, ensuring a seamless shopping experience for millions of customers.”

ADITYA TIWARI

SDE III, Data Platform, Flipkart



Industry

Retail and CPG

Country

Austria

Google Cloud database

[Spanner](#)

Additional Google

Cloud products

Google Kubernetes Engine

REWE Group doubles transaction capacity with Spanner

Challenge

As one of Europe's largest retail and tourism groups, REWE Group needed a scalable, high-availability solution to handle traffic spikes from millions of customers across online and in-store applications – without disrupting operations.

Solution

REWE Group deployed Spanner within a Kubernetes cluster on Google Cloud. Spanner's 99.999% availability enabled a new transactional ledger system that improved speed and reliability for customer purchases. Fully managed services allowed developers to focus on innovation, while Google Cloud's sustainability initiatives supported REWE's environmental goals.

Outcomes

- Soon after launch, REWE Group processed nearly **500 million successful transactions**
- The system **doubled transaction processing capacity** with real-time updates
- **Faster API calls** in point-of-sale applications improved the customer experience

[Read the full story](#)

“Query latency is always a critical thing for us because we are deeply integrated into the point-of-sale applications in our store. If applications are too slow, it compromises the customer experience. However, thanks to Spanner, we are able to complete API calls extremely fast.”

ANDREAS RÖHRENBACHER

Chief Product Owner, IT, REWE Group

**Industry**

Retail and CPG

Country

Germany

Google Cloud database[AlloyDB for PostgreSQL](#)**Additional Google****Cloud products**

Vertex AI

Tchibo brews up 10x faster customer insights with AlloyDB for PostgreSQL

Challenge

Coffee retailer and lifestyle brand Tchibo had reached the limits of its cloud database, which couldn't scale to meet the complexity of its customer feedback data. This made the process of generating insights extremely labor-intensive and delayed decision-making.

Solution

Tchibo migrated to AlloyDB for PostgreSQL to modernize analytics and gain AI-driven insights. With high-performance processing and built-in vector search, Tchibo now extracts insights in real time, making customer feedback more actionable. Automating report generation has also freed employees from manual work, which allows for faster decision-making across product offerings and sales channels.

Outcomes

- AlloyDB for PostgreSQL's high performance analytics and RAG workflows delivered **10x faster queries**
- **Instant report generation** replaced days of manual effort, so teams can respond to feedback faster
- Fully managed AlloyDB for PostgreSQL **reduced operational overhead** by scaling effortlessly as data volumes grow

[Read the full story](#)

“AlloyDB for PostgreSQL provided a powerful solution to the limitations we faced with our old database. Its advanced analytics capabilities, built-in vector search, and familiar PostgreSQL foundation offered the speed, adaptability, and usability we needed to serve up insights as fresh and fast as our coffee.”

HENNING KOSMALLA

Principal Data Scientist, GenAI, Tchibo

Financial services

Strengthen every transaction, every trade, and every claim with a database built for security, scale, and speed.



Deutsche Bank

**Industry**

Financial services

Country

Germany

Google Cloud database[Spanner](#)

Deutsche Bank migrates to Spanner to enhance scalability

Challenge

Deutsche Bank needed a modern, scalable, and resilient database to support its online banking platform and handle millions of transactions per day. Its existing system required significant operational overhead and was reaching its scalability limits at a time when ensuring high availability and real-time data consistency was critical.

Solution

By adopting Spanner, Deutsche Bank successfully modernized its core banking infrastructure, consolidating massive datasets and enabling real-time, high-volume transaction processing. Spanner's fully managed, globally distributed architecture provided the scalability and resilience needed to support Deutsche Bank's online banking operations, while reducing latency and operational complexity.

Outcomes

- Distributed architecture ensures **consistency across 19 million product contracts and 12 million customer records**
- 99.999% availability guarantees **uninterrupted banking services**, even during peak transaction periods
- The **migration was completed in a single weekend**, with no impact on customer transactions or access

[Read the full story](#)

“Scaling in high-availability environments can be challenging, but Spanner does all the heavy lifting. Spanner scales infinitely and allows Deutsche Bank to start small and easily scale up and down as needed.”

MICHAEL OTMAR KAISER

Lead Engineer, Deutsche Bank AG



Industry

Financial services

Country

United States

Google Cloud database

[Bigtable](#)

Additional Google

Cloud products

BigQuery, Dataproc, Dataflow,
Google Kubernetes Engine

Equifax uses Bigtable to reinvent itself through technology

Challenge

On a mission to transform its services, Equifax wanted to build a foundation that would allow it to close 57 data centers, break down data silos, ingest new data sources, and ensure the safety, security, and governance of all its data. Hoping that this new data fabric would accelerate innovation, Equifax aimed to get new ideas into production quickly while improving collaboration and insights.

Solution

Equifax used Bigtable as a key data architecture component for its data fabric. As a fully managed service, Bigtable helps Equifax increase the speed and scale of innovation by rapidly ingesting data from suppliers, capturing and organizing the data, and serving it to users so they can build new products.

Outcomes

- Equifax gained **instantaneous access to credit data**, something only a cloud-native bureau can do
- **Rapid responses to identity questions** with keying and linking across a repository of 3 billion credit observations
- The **highest level of security, data protection, and governance** to meet regulations

[Read the full story](#)

“Moving to Google Cloud has enabled Equifax to accelerate how we bring innovation into the market.”

CECILIA MAO

Chief Product Officer, Equifax



Industry

Financial services

Country

Brazil

Google Cloud database

[Spanner](#)

Additional Google

Cloud products

Apigee API Management,
BigQuery, Google Kubernetes
Engine

CERC powers 100,000 transactions per second with Spanner

Challenge

As a financial market infrastructure (FMI), CERC needed to build a scalable, digital-first platform to register and settle receivables as credit collateral – while supporting high transaction volumes and meeting Brazil’s Central Bank regulations.

Solution

CERC built its infrastructure on Google Cloud to support rapid growth and high transaction volumes. Spanner delivered the scalability and availability needed to process 100,000 transactions per second. In addition, Google Kubernetes Engine supported its microservices-based architecture, BigQuery enabled data sharing with financial agents, and Apigee laid the foundation for an API-first approach.

Outcomes

- Scalable infrastructure now supports over **100,000 transactions per second**
- A mandated Central Bank project, stagnant for over a decade, was **successfully launched into production in under six months**
- **Complex environments are now easily managed** with scalable, cloud-native infrastructure

[Read the full story](#)

“The ability to add more resources under a transparent model is the most relevant aspect for us. It’s a performance-related factor that provides added value. We want to make the most of the tools while making room for expansion and internal learning.”

MARCELO MAZIERO

Co-Founder and Chairman, CERC

current

Industry

Financial services

Country

United States

Google Cloud database[Spanner](#)

Current uses Spanner to build a resilient banking platform

Challenge

Current utilizes an innovative approach to quickly provide financial solutions that are precisely tailored to members' immediate needs and are inherently designed to adapt as those needs evolve. To create a scalable and robust technological foundation for financial services, Current needed to build a modern, core banking system to power its platform.

Solution

Current migrated from a self-hosted graph database to Spanner, gaining the consistency, scalability, and reliability it needed to support its core banking platform. Leveraging Spanner to power its graph service, Current can deliver reliable and efficient financial services, which are critical for building and maintaining member trust.

Outcomes

- Current **reduced recovery time objective and recovery point objective by more than 10x** – cutting times to just one hour
- Zero availability-related incidents **enhance customer retention and improve the developer experience**
- Current can easily scale to nearly **5,000 transactions per second** with Spanner while reducing the team's maintenance burden

[Read the full story](#)

“Spanner emerged as the ideal solution, fulfilling all our requirements. It offers consistent writes, horizontal scalability, and the ability to maintain low read latency even under a heavy load. Its seamless scalability – particularly the decoupling of compute and storage resources – proved invaluable in adapting to our dynamic consumer environment.”

TREVOR MARSHALL

CTO and Co-Founder, Current

**Industry**

Financial services

Country

Singapore

Google Cloud databases[AlloyDB for PostgreSQL](#),[Memorystore](#)**Additional Google****Cloud products**

Datastream, Database

Migration Service, BigQuery,

Google Kubernetes Engine

Galxe migrates to AlloyDB for PostgreSQL, cutting costs by 40%

Challenge

Galxe is a leader in Web3 identity and engagement infrastructure. The company needed a more cost-efficient and scalable database to support its flagship product, Galxe Quest, which serves over 26 million users. Its previous solution, Amazon Aurora, struggled with the scale and cost of growing read and write operations.

Solution

Galxe migrated from Amazon Aurora to AlloyDB for PostgreSQL to support its expanding user base and simplify data access for developers. The team used Google Cloud's Database Migration Service to replicate data continuously with minimal downtime. AlloyDB's PostgreSQL compatibility enabled a smooth transition and now serves as Galxe's single source of truth for on-chain and off-chain data.

Outcomes

- Galxe **reduced database costs by 40%** after migrating to AlloyDB
- The platform now **supports over 26 million users with seamless scalability**
- **Near-zero downtime during migration** helped maintain platform availability

[Read the full story](#)

“We trust AlloyDB for PostgreSQL for its flexibility, near-zero downtime, and superior performance. Acting as our single source of truth, AlloyDB for PostgreSQL securely houses millions of on-chain and off-chain user data records. This enables our developers to access granular datasets for building their loyalty programs on the blockchain.”

ZHONGTIAN WANG

Head of Infrastructure, Galxe



Industry
Financial services

Country
Norway

Google Cloud database
[Cloud SQL for PostgreSQL](#)

Additional Google Cloud products
Google Kubernetes Engine,
Vertex AI

Deckmatch powers insights for venture capitalists using Cloud SQL for PostgreSQL

Challenge

To improve the efficiency and strategic focus of venture capitalists, financial services company Deckmatch needed a powerful and flexible database to handle large volumes of data, support advanced vector-based searches, and deploy AI to deliver deep insights into potential investment opportunities.

Solution

Using Cloud SQL for PostgreSQL with pgvector, Deckmatch enhanced investment decisions with fast, contextual search capabilities and comprehensive startup intelligence, freeing venture capitalists from manual data processing.

Outcomes

- **Pgvector enables fast and efficient similarity searches**, supporting advanced embedding algorithms for competitive mapping
- Using **a fully managed service eliminates operational overhead** like backups and updates
- Processing extensive company data empowers venture capitalists to **make timely and informed investment decisions**

[Read the full story](#)

“Cloud SQL’s intuitive interface and seamless integration with other Google Cloud Platform tools makes setup and ongoing management incredibly straightforward... Plus, because Cloud SQL enables us to flexibly scale database resources on demand, it grows as our application does.”

WALID MUSTAPHA
Co-Founder and CTO, Deckmatch



Industry
Financial services

Country
United States

Google Cloud database
[AlloyDB for PostgreSQL](#)

Additional Google Cloud products
Google Kubernetes Engine,
Pub/Sub

AlloyDB for PostgreSQL fuels 50% faster margin calculations for Apex Fintech

Challenge

Apex Fintech Solutions Inc. enables modern investing and offers wealth management tools through an ecosystem of frictionless platforms, APIs, and services. Faced with a rapidly growing customer base, Apex sought to migrate to the cloud so it could deliver more timely and accurate margin calculations.

Solution

Apex chose to migrate to Alloy DB for PostgreSQL because it met its high availability and disaster recovery requirements right out of the box. The new solution allows it to use real-time data from APIs to calculate margin on demand and determine risk in seconds.

Outcomes

- Apex **reduced processing time by 50%**, enabling margin calculations for 100,000 accounts in just one minute
- Flexible architecture allowed Apex to **expand existing service offerings and increase customer value**
- AlloyDB for PostgreSQL's columnar engine eliminated the need for an additional analytical system, resulting in **lower overall costs**

[Read the full story](#)

“Total system availability and performance were critical considerations for us as a highly regulated financial institution. Because AlloyDB for PostgreSQL was compatible with our high availability and disaster recovery requirements right out of the box, we quickly adopted this solution with confidence.”

ANTRO PETER

Senior Director of Technology, Apex Fintech Solutions



FLUIDEFI

Industry
Financial services

Country
United States

Google Cloud database
[AlloyDB for PostgreSQL](#)

Additional Google Cloud products
Database Migration Service

FLUIDEFI reduces costs by 60% and increases response times by 3x with AlloyDB for PostgreSQL

Challenge

FLUIDEFI is a SaaS platform for institutional investors in the decentralized finance space, providing risk qualification, portfolio modeling, and other services to plug these gaps. Given the amount of data it manages in near real-time – over 2 billion records – it requires a robust managed database service.

Solution

AlloyDB for PostgreSQL powers several critical services for FLUIDEFI in near real-time, generating accurate financial data, actionable insights, risk ratings, and audit trails. For FLUIDEFI’s end users, these services streamline risk qualification, portfolio model creation, auditable financial reporting, investment tracking, and investment strategy testing.

Outcomes

- **Increased its platform’s response speed by 3x** with capacity to support more transactions per second
- **Reduced costs by 60%** for the same compute capacity as it used previously
- Improved **scalability** without worrying about decreased data delivery speeds

[Read the full story](#)

“We see AlloyDB for PostgreSQL as an essential tool to achieve our vision – to capture a significant share of the rapidly growing DeFi market.”

LOUIS SIRICO
CTO and Co-Founder, FLUIDEFI

**Industry**

Financial services

Country

United States

Google Cloud database[Bigtable](#)**Additional Google****Cloud products**

Dataflow, Dataproc, Google

Cloud Storage

Symphony manages millions of messages with Bigtable, saving 40% on database costs

Challenge

As a pioneer in secure communication and collaboration, Symphony fosters information exchange for financial institutions worldwide. Symphony needed a managed database that could handle large amounts of data while also providing the scalability, reliability, and flexibility required to support its growing user base.

Solution

Symphony migrated to Bigtable from a self-managed, open-source database that was time consuming and expensive to maintain. The move to Bigtable reduced operational costs and time spent on maintenance, freeing up the team to focus on automation, scalability, and other critical tasks.

Outcomes

- Symphony **reduced operational costs by 40%** and decreased administration and maintenance activities
- Low latency and high stability allow Symphony to **access messages of varying sizes in 15ms (p95) or less**
- Seamless data migration to production cluster in minutes ensures **minimal customer downtime and zero impact on production systems**

[Read the full story](#)

“Moving forward, this Bigtable migration will continue to unlock new possibilities. Storing massive amounts of data at scale becomes possible without incurring performance penalties. Global replication offers new opportunities for scalability and resilience.”

OLIVIER RICHAUD

VP, Platforms and SRE, Symphony

sanitas

Industry

Financial services

Country

Switzerland

Google Cloud database[Cloud SQL for PostgreSQL](#)**Additional Google****Cloud products**

Google Kubernetes Engine

Sanitas modernizes its database with a true DevOps model

Challenge

Sanitas, a Swiss healthcare insurer that serves over 800,000 customers, needed to modernize its customer-facing infrastructure. The goal: shift from on-premises systems to a cloud-first model that could support faster innovation and improve developer efficiency.

Solution

Sanitas adopted Cloud SQL for PostgreSQL to replace legacy Oracle databases and support a DevOps-friendly infrastructure. The team migrated its CRM and customer service portal, adapting code and coordinating with vendors to preserve performance. Running Cloud SQL with GKE-based microservices improved reliability and limited the impact of incidents. A phased migration approach and built-in insights helped the team minimize downtime and shift focus to development.

Outcomes

- A data isolation microservice pattern provides **high reliability during peak loads**
- A flexible self-service provisioning model and out-of-the-box diagnosis capabilities **improved developer efficiency**
- **Increased bandwidth** enables operations and data teams to fine-tune workloads and assist application development teams

[Read the full story](#)


“The transition to Cloud SQL not only enhanced the ability to isolate microservices, but also propelled us towards a genuine DevOps operational model.”

MARCEL AMSLER

DevOps Technical Lead, Sanitas

Technology

Whether you're building APIs, platforms, or software, Google Cloud databases keep your backend running strong.



**Industry**

Technology

Country

United States

Google Cloud database[Spanner](#)**Additional Google****Cloud products**

Cassandra to Spanner Proxy

Adapter

Yahoo modernizes its database infrastructure with Spanner

Challenge

Yahoo needed a scalable, reliable, and cost-effective database solution for its Contacts workload. After deciding to transition from Apache Cassandra to Spanner, the team wanted to avoid rewriting CQL queries and minimize engineering overhead.

Solution

Using the Cassandra to Spanner Proxy Adapter, Yahoo seamlessly adopted Spanner without reworking its application logic. The adapter allowed for a smooth transition, freeing engineering teams to focus on higher-priority initiatives, while gaining the benefits of Spanner's global availability and built-in redundancy.

Outcomes

- Yahoo transitioned to Spanner while **keeping existing CQL queries unchanged**
- Cassandra Adapter enabled a **flexible migration strategy**, allowing engineering teams to focus on innovation
- Spanner provides Yahoo with the **scale, redundancy, and cost-effectiveness** needed for a business of its size

[Read the full story](#)

“Our migration strategy has more flexibility, and we can focus on other engineering activities while utilizing the scale, redundancy, and support of Spanner without updating the codebase.”

PATRICK JD NEWMAN

Principal Product Manager, Core Mail and Analytics, Yahoo



Industry
Technology

Country
United States

Google Cloud database
[Cloud SQL](#)

Additional Google Cloud products
Database Migration Service,
Google Kubernetes Engine

Google Nest seamlessly migrates terabytes of data to Cloud SQL

Challenge

Google Nest, known for its innovative smart home products, housed terabytes of subscription data in legacy infrastructure on AWS. These workloads were self-managed services, so they required extensive site reliability engineering resources for 24/7 monitoring.

Solution

Using Google Cloud's Database Migration Service, Google Nest modernized its infrastructure by migrating multiple MySQL databases and terabytes of data to Cloud SQL – with near-zero downtime and no data loss. The fully managed, automated database service provided high availability and seamless performance for critical subscription services.

Outcomes

- Cloud SQL **cut operational costs** by adding storage capacity and reducing maintenance overhead
- A managed service approach **freed up at least 10% of bandwidth** for the site reliability engineering team
- Unifying legacy subscription data in Google Cloud **reduced p50 latency by 25%**

[Read the full story](#)

“As a secure, fully managed, and automated service, Cloud SQL allowed us to hand off database management, cut our operational costs, and maintain the reliability and performance we needed for critical subscription services.”

ARPIT GOYAL
Tech Lead, Google Nest



Industry
Technology

Country
United States

Google Cloud databases
[Bigtable](#), [Spanner](#)

Sabre supports over 1 billion travelers with Bigtable and Spanner

Challenge

Sabre, a travel software provider, needed a low-latency database to power more than 12 billion annual shopping requests for over 1 billion travelers. It needed a solution that could quickly deliver results across mobile apps, through third-party travel sites, and directly with airlines.

Solution

Sabre selected Spanner to manage airline reservations, prioritizing global consistency, ACID compliance, and high reliability across multiple availability zones. Spanner underpins critical airline operations, including large-scale rebooking events. For flight shopping, Sabre implemented Bigtable as a low-latency NoSQL database that serves large volumes of shopping results cost effectively and with fast response times.

Outcomes

- Bigtable's **single-digit millisecond response time for multi-petabyte tables** serves results with 99.999% availability
- Spanner processes over **1 billion requests per second** at peak
- The ability to cache high volumes of shopping results **reduces the cost of compute usage**

[Read the full story](#)

“In an industry as demanding as travel, accelerating our most critical applications using technology unique to Google Cloud means less time spent optimizing latency and consistency, and more time spent innovating.”

ANDREW GASPAROVIC

Vice President and Chief Architect, Sabre Labs



Industry
Technology

Country
United Kingdom

Google Cloud database
[Cloud SQL](#)

Additional Google Cloud products
BigQuery, Google Kubernetes Engine

Auto Trader increases release cadence by 140% with Cloud SQL

Challenge

Auto Trader, the UK's leading online car marketplace, relied on a large, monolithic database that limited agility. Migrating to the cloud was critical to accelerating feature delivery.

Solution

Auto Trader used Google Cloud Database Migration Service to facilitate its migration to Cloud SQL, a fully managed relational database service. Post-migration, Cloud SQL supported critical services like its Vehicle Data Service and inventory management with strong performance improvements and enhanced scalability.

Outcomes

- **Release frequency increased to over 140% year over year**, peaking at 458 releases in a single day
- In one year, Cloud SQL and Google Kubernetes Engine **supported 36,000 releases with a success rate of 99.87%**
- The new solution can **dynamically scale resources for critical services** within a 5-minute window

[Read the full story](#)

“Cloud SQL’s fully managed services took away the headache of database maintenance that would typically take up a lot of our energy.”

MOHSIN PATEL

Principal Database Engineer, Auto Trader

**Industry**

Technology

Country

Japan

Google Cloud databases[Bigtable](#), [AlloyDB for](#)[PostgreSQL](#)**Additional Google****Cloud products**

BigQuery

PLAID puts the “real” in real-time user analytics with Bigtable

Challenge

PLAID, which provides connections to more than 11,000 US banks and credit unions, needed to re-architect its core customer experience platform, Blitz. In order to handle over 100K events per second with millisecond processing times, Blitz required a database designed for high scalability and low latency.

Solution

To achieve high scalability and low latency for its real-time analytics engine, PLAID adopted Bigtable. Leveraging Bigtable’s low-latency key-value store and range scan capabilities, PLAID built a horizontally scalable, distributed queue that maintains latency within 10ms.

Outcomes

- PLAID revamped its core analytics engine with Bigtable to achieve **consistent, real-time analytics when traffic is high**
- PLAID can operate its real-time distributed queue at **less than half the cost** compared to running the workload on alternative services
- Bigtable’s **auto-scaling features reduce operational cost** and allow for greater flexibility and less management

[Read the full story](#)

“By leveraging the power of Bigtable, we believe that businesses can unlock new levels of performance and consistency in their real-time analytics engines, ultimately leading to better user experiences and more insightful decision-making.”

JUN KUSAHANA

Senior Software Engineer, PLAID



Industry
Technology

Country
Japan

Google Cloud databases
[Cloud SQL for SQL Server](#),
[Cloud SQL for PostgreSQL](#),
[Memorystore for Redis](#)

Visual Research reduces costs by 35% and improves query processing by up to 20% with Cloud SQL

Challenge

Visual Research helps real estate businesses and their customers execute rental and management processes. As it transitioned to a more digital experience, Visual Research faced challenges with increasing license costs, inefficient resource allocation, and keeping up with manual maintenance tasks.

Solution

Visual Research turned to Google Cloud for scalable, flexible, high-speed databases. Cloud SQL for SQL Server and Memorystore for Redis support rental management systems, while Cloud SQL for PostgreSQL and Memorystore for Redis support brokerage and sales. These fully managed services help agencies handle seasonal demand and make faster, data-driven decisions.

Outcomes

- High availability and streamlined operations **reduced costs by 35%**
- A reduction in manual tasks and storage management **saved 4–8 hours per month**
- Faster data access **improved query performance up to 20%**

[Read the full story](#)

“The benefits of Cloud SQL for SQL Server are particularly significant for our team, as we lack platform engineers. This transition allows us to place a more substantial focus on our core objective – application development.”

SHUN WATANABE

Chief Technology Officer, Visual Research

**Industry**

Technology

Country

United States

Google Cloud database[Firestore](#)**Additional Google Cloud products**

Google Kubernetes Engine, Pub/Sub, Vertex AI

HighLevel builds an AI marketing platform with Firestore that scales to over 30 billion documents

Challenge

HighLevel, a fast-growing SaaS platform for marketing agencies, needed a database that could support unpredictable write spikes from a few hundred to hundreds of thousands of requests per second. Its previous cloud document database required manual provisioning and frequent sharding, which slowed releases.

Solution

In need of a database solution that could seamlessly scale and handle demanding write requirements, HighLevel migrated to Firestore. Its serverless architecture, point-in-time recovery, and scheduled backups have enabled HighLevel to increase productivity, scalability, reliability, and more.

Outcomes

- Firestore has **boosted developer productivity by 55%**, allowing HighLevel to focus on product innovation
- Firestore handles workloads with spikes of **up to 250K requests per second and 5M real-time queries**
- **Real-time sync capabilities power real-time dashboards** without the need for complex socket infrastructure

[Read the full story](#)

“Firestore has been instrumental in our ability to scale rapidly, enhance developer productivity, and deliver innovative AI-powered solutions. We are confident that Firestore will continue to be a cornerstone of our technology stack as we continue to grow and evolve.”

KARAN AGARWALDirector of Engineering – CRM, AI & Platform,
HighLevel



Industry
Technology

Country
United States

Google Cloud database
[Cloud SQL](#)

Additional Google Cloud products
BigQuery, Vertex AI, Google Kubernetes Engine

Cart.com builds a unified analytics solution with Cloud SQL

Challenge

As ecommerce surged during the COVID-19 pandemic, brands were pushed to rethink how they operate online. Meeting evolving customer expectations required creative, forward-thinking strategies – but many were stuck using disconnected tools for marketing, fulfillment, and analytics that made it hard to move quickly or make informed decisions. Cart.com recognized these challenges early and was ready to help brands adapt.

Solution

Cart.com built a unified ecommerce platform on Google Cloud that centralizes data across marketing, fulfillment, and storefronts. The platform processes 40 million real-time events daily, using Cloud SQL and BigQuery for data storage and analytics, Vertex AI for machine learning and AI-driven insights, and Google Kubernetes Engine for scalable application deployment.

Outcomes

- Full load tests with 2,500 simultaneous virtual users resulted in **zero scaling errors**
- Retailers using Cart Unified Analytics see **margins increase by up to 20%**
- Brands can deploy Cart Unified Analytics and **start using it within 60 minutes**

[Read the full story](#)

“We’ve built it the right way from the beginning, working with the best minds at Google Cloud. We always take a technology first mindset in M&A and this enables the privilege to think blue sky and build it as we go.”

CHASE ZIEMAN

Chief Data Science Officer, Cart.com

**Industry**

Technology

Country

United States

Google Cloud databases[AlloyDB for PostgreSQL](#),[Firestore](#)**Additional Google****Cloud products**

BigQuery, Database

Migration Service, Elastic on

Google Cloud, Pub/Sub

Endear averages a 6x increase in transactions per second with AlloyDB for PostgreSQL

Challenge

Retail and ecommerce software provider, Endear, struggled with database scalability and performance issues as its customer base grew. Its existing database service couldn't handle the increasing volume of data integration and real-time insights required for the company's customer relationship management (CRM) platform.

Solution

Endear migrated to AlloyDB for PostgreSQL to support its growing transaction volume and complex data architecture. AlloyDB's PostgreSQL compatibility made the transition smooth, while its high performance met the demands of real-time CRM workloads. During migration, Google Cloud's Database Migration Service enabled continuous syncing with zero downtime. Firestore also supports real-time app updates across the platform, enhancing responsiveness for end users.

Outcomes

- Endear's AlloyDB for PostgreSQL cluster **serves 6x the number of connections** compared to its previous solution
- **Transaction volume rose 6x**, reaching 10K per second on the read cluster and 5K per second on the primary
- AlloyDB for PostgreSQL can sustain and scale workloads easily, with a **P99 aggregated query latency under 10ms**

[Read the full story](#)

“The strategic use of managed services from Google Cloud enables us to focus on delivering value to our customers without having to handle the complexities of data management and scaling.”

JP GRACE

CTO, Endear



Industry
Technology

Country
Brazil

Google Cloud databases
[AlloyDB for PostgreSQL](#),
[Firestore](#)

**Additional Google
Cloud products**
Cloud Run

B4A improves ecommerce platform query times by 90% with AlloyDB for PostgreSQL

Challenge

As adoption grew for Brazilian beauty tech startup, B4A, its Microsoft SQL Server database struggled to efficiently process queries for its catalog of 10,000+ beauty products. Some product collections had load times of up to 12 seconds, which slowed the shopping experience and limited scaling capabilities.

Solution

B4A migrated its backend to AlloyDB for PostgreSQL to optimize query performance and gain seamless scalability. The fully managed database simplified maintenance, eliminated infrastructure management overhead, and integrated smoothly with the company's existing Google Cloud environment. B4A also uses Firestore to support real-time use cases like its micro-influencer monetization platform, benefiting from serverless triggers and easy integration with Cloud Functions.

Outcomes

- **Load times for catalog queries improved by up to 90%**, from 12 seconds to as low as 0.25 seconds
- B4A **reduced costs by only paying for what it uses** with AlloyDB for PostgreSQL's predictable and transparent pricing
- **Database management became more efficient**, allowing developers to work on high-value projects

[Read the full story](#)

“In a startup like ours, resources are limited, so the ease of use and integration and lack of infrastructure requirements allow for a much better overall result. Now, B4A developers can handle almost everything without support from the operations team.”

JAN REIHLE
CEO and Founder, B4A



Industry
Technology

Country
United States

Google Cloud database
[Memorystore for Redis Cluster](#)

Unity Ads scales to 10M operations per second with Memorystore

Challenge

Unity Ads, a mobile advertising platform, previously relied on a DIY Redis setup that was difficult to scale and maintain. Performance instability during spikes and Kubernetes upgrades made it hard to meet the demands of real-time ad serving.

Solution

To eliminate these bottlenecks, Unity Ads migrated key workloads – including session data, valuation cache, and distributed locks – to Memorystore for Redis Cluster. The team used double-writes to transition with no service disruption and completed major migrations in as little as 15 minutes. With one-click persistence and built-in high availability, Memorystore gave Unity a simpler, more reliable foundation for short-lived and high-throughput data use cases.

Outcomes

- Unity's infrastructure now **supports up to 10 million Redis operations per second for a single instance**
- System stability improved significantly, with **consistent low latency** even during scaling events
- **Operational overhead dropped** as engineers no longer manage upgrades, tuning, or Redis internals

[Read the full story](#)


“With Memorystore, we’re now getting a fully managed and much more reliable (99.99% SLA) and scalable solution for a similar cost as our previous self-managed Redis deployment.”

EREN BOZ

Senior Software Engineer, Unity

Startups and digital natives

From one user to virtually limitless, scale seamlessly with data that's always fast and reliable.



Uber

Industry

Startups and digital natives

Country

United States

Google Cloud database

[Spanner](#)

Additional Google

Cloud products

Google Kubernetes Engine

Uber scales its fulfillment platform globally with Spanner

Challenge

Uber's fulfillment platform powers billions of database transactions a day across trips, deliveries, and logistics. As demand grew and new verticals launched, Uber needed a globally consistent database that could scale horizontally without increasing operational overhead.

Solution

Uber migrated its core platform workloads from a legacy NoSQL architecture to Spanner. The new solution eliminated complex sharding logic, reduced latency through a hybrid cloud network design, and maintained referential integrity at the application layer. Spanner's external consistency allowed Uber to simplify distributed transactions and avoid manual recovery workflows, while built-in observability and autoscaling features helped maintain performance and control costs.

Outcomes

- Spanner's multi-region deployment **provided 99.999% availability for mission-critical services**
- Spanner **reduced infrastructure complexity** to boost developer productivity
- Custom autoscaling and on-prem caching **cut operational costs** tied to peak traffic loads

[Read the full story](#)

“We’ve solved deep technical issues in our network routing infrastructure, productionized our application workload against Spanner infrastructure, and continuously optimized the architecture to improve the latency, availability, and database performance.”

ANKIT SRIVASTAVA

Senior Staff Engineer, Uber



Industry

Startups and digital natives

Country

United States

Google Cloud databases

[Spanner](#), [Cloud SQL](#)

Additional Google

Cloud products

Compute Engine

Vimeo builds a fully responsive video platform with Spanner

Challenge

As a leading video software solution powering billions of views, Vimeo needed to modernize its infrastructure to scale globally and deliver a seamless playback experience for both everyday uploads and high-demand content.

Solution

To improve availability and responsiveness, Vimeo re-architected its platform with Spanner, gaining global scale and relational semantics with minimal overhead. Spanner's point-in-time recovery simplified data validation during migration. Vimeo also uses Cloud SQL for lower-volume services that don't require Spanner's scale or availability.

Outcomes

- Spanner supported **seamless global video playback** with 99.999% availability
- Vimeo **scaled to 50.8 billion rows and 4.5TB of storage** across 16 Spanner nodes
- The platform **supports dynamic playback and metadata search across billions of videos**

[Read the full story](#)

“With Spanner, we found all the benefits of a relational semantics database with global scale, allowing us to add more nodes with the push of a button.”

SERGIO SALVATORE

Senior Director, Engineering, Vimeo

**Industry**

Startups and digital natives

Country

United States

Google Cloud database[Bigtable](#)**Additional Google
Cloud products**

BigQuery, Dataproc

Box migrates critical storage to Bigtable, decreasing its footprint by 85%

Challenge

Box is a recognized leader in cloud-based content management, collaboration, and file sharing. The company needed a scalable database solution to decrease the operational burden of managing three on-premises Apache HBase clusters.

Solution

Box migrated 600 TB of data from its three HBase clusters to Bigtable to reduce operational maintenance, improve scalability, and lower costs. Bigtable autoscaling now allows Box to run lighter secondary clusters and scale up only when needed.

Outcomes

- Box **reduced costs and operational maintenance with an 85% smaller storage footprint** by leveraging autoscaling features
- Bigtable autoscaling enabled Box to speed up development, **completing MapReduce jobs in under 24 hours that previously took days**
- Combining Bigtable with Google services like BigQuery, Box **gained faster data analysis tools**

[Read the full story](#)

“Our team now has much less overhead related to managing our database. In the past, we would constantly have to move around HBase traffic to perform security patches. Now, we don’t need to worry about managing that at all.”

MINDY YANG

Senior Software Engineer, Box

**Industry**

Startups and digital natives

Country

United States

Google Cloud database[Bigtable](#)**Additional Google
Cloud products**

BigQuery, Cloud Armor,
Cloud Build, Compute Engine,
Dataflow, Google Cloud
Storage, Google Kubernetes
Engine, Looker

[Watch the full story](#)

Bitly's move to Bigtable increased resiliency, speed, and scalability

Challenge

Bitly, the link and QR code management platform, was outgrowing its self-managed MySQL database. With data for over 40 billion links and counting, basic operational tasks like software and security upgrades became substantial challenges. Additionally, the backup and restore process was costly and time-consuming, leading Bitly to search for a new solution.

Solution

Bitly migrated 80 billion rows of its core link data from a self-managed MySQL database to Bigtable. The move to Bigtable provided greater resiliency, speed, and scalability for future growth.

Outcomes

- **Export speeds increased to 15 million rows per second** and restore speeds neared 2 million writes per second
- Bitly **reduced the amount of data it stores by nearly 50%** by migrating to Bigtable, freeing up significant storage space
- Bitly **gained resiliency and redundancy** with Bigtable backups, multi-region replication, and geo-distribution

“Not only did Bigtable meet our technical requirements and operational needs, but it also sets us up for future growth. Its ability to scale seamlessly over time while improving our system availability SLA was a major factor in our decision.”

ZOE MCCORMICK

Senior Software Engineer, Bitly

character.ai

Industry

Startups and digital natives

Country

United States

Google Cloud databases

[AlloyDB for PostgreSQL](#),

[Memorystore for Redis Cluster](#),

[Spanner](#)

Character.AI powers real-time conversations with low-latency, scalable databases

Challenge

Character.AI, an AI chatbot company with over 20 million monthly users, needed a database infrastructure that could support rapid growth without compromising latency, scalability, or developer productivity.

Solution

To support different layers of its stack, Character.AI chose a combination of Google Cloud databases tailored to each use case. Spanner handles the fast-growing chat backend with high availability and virtually unlimited scale. AlloyDB for PostgreSQL powers its frontend chat and system of engagement with consistent low-latency reads. For caching, the team migrated to Memorystore for Redis Cluster, eliminating the complexity of proxy-based sharding while maintaining high cache hit rates and low tail latencies.

Outcomes

- Spanner **ingests terabytes of data daily** with 99.999% availability
- Memorystore **reduced p99 latency by 66%** for high-traffic endpoints
- AlloyDB for PostgreSQL **scaled read capacity by 20x** with no app code changes

[Read the full story](#)

“We no longer have to worry about manually maintaining proxies or hashrings... Google Cloud allows us to focus on our core technologies and user experience.”

JAMES GROENEVELD

Research Engineer, Character.AI

**Industry**

Startups and digital natives

Country

United States

Google Cloud databases[Cloud SQL for PostgreSQL](#),[Memorystore for Redis](#)

Linear scales vector search with pgvector on Cloud SQL

Challenge

Since 2019, Linear has supported global product development workflows for businesses through its project and issue-tracking system. For new production workloads storing large vectors, it needed to upgrade its databases to a cost-efficient, scalable, and reliable solution with strong vector search support.

Solution

Linear adopted Cloud SQL for PostgreSQL with pgvector to bring semantic similarity search into production. Additionally, it uses Google's managed Memorystore for Redis as an event bus and cache. The team migrated a large development dataset to production by partitioning the issues table and indexing each segment. Cloud SQL's managed service eliminated the need for a dedicated operations team while supporting a growing data footprint in tens of terabytes.

Outcomes

- By utilizing Cloud SQL, Linear **scaled into tens of terabytes** without additional engineering overhead
- Pgvector-powered similarity search **improved bug reporting accuracy and customer support workflows**
- Fully managed infrastructure freed developers to **focus on product improvements**

[Read the full story](#)

“Cloud SQL for PostgreSQL has proven invaluable for Linear. Because we do not have a dedicated operations team, relying on managed services is crucial.”

TOM MOOR

Head of US Engineering, Linear

**Industry**

Startups and digital natives

Country

Israel

Google Cloud database[Cloud SQL for PostgreSQL](#)**Additional Google****Cloud products**

BigQuery, Dataflow

Lightricks boosts search retrieval rates by 40% with pgvector support in Cloud SQL

Challenge

Lightricks, the company behind creative apps like Videoleap, needed more responsive and intuitive search capabilities to help users navigate its vast library of video templates. The existing keyword-based search was too rigid, making it difficult to deliver relevant results based on user intent.

Solution

Lightricks adopted the pgvector extension in Cloud SQL for PostgreSQL to bring semantic search to Videoleap. This allowed it to match queries to templates using vector embeddings. Its team chose Cloud SQL to avoid syncing issues, reduce learning curves, and streamline deployment with the infrastructure it already had.

Outcomes

- Implementing pgvector in Cloud SQL led to a **40% increase in number of retrievals and template usage from retrieved results**
- **Response times (p90) for search queries decreased** from 1–4 seconds to under 100 milliseconds
- Enhanced search functionality enabled **querying millions of embeddings with high accuracy**

[Read the full story](#)

“When Cloud SQL for PostgreSQL rolled out pgvector support, we knew it was the right choice.... Its streamlined approach reduces development overhead and minimizes the risk of data inconsistencies, making it a more efficient and reliable solution for many use cases.”

DAVID GANG

Tech Lead, Lightricks

WRITER

Industry

Startups and digital natives

Country

United States

Google Cloud databases

[AlloyDB for PostgreSQL](#),
[Bigtable](#), [Memorystore](#)

Additional Google

Cloud products

Datastream for BigQuery,
Google Kubernetes Engine

Writer.com powers generative AI at scale with Google Cloud databases

Challenge

Writer.com, an enterprise-grade generative AI platform, needed to scale its database infrastructure to support explosive growth – from 128 million parameter LLMs to trillions of API calls per month. Managing MySQL at this scale would require too many resources and limit the team’s ability to focus on customer needs.

Solution

Writer.com migrated to a mix of Google Cloud databases, anchored by AlloyDB for PostgreSQL, to support around 100 microservices running on GKE. The team used AlloyDB’s full text search for microservice performance, and began exploring AlloyDB AI for semantic search. Bigtable and Memorystore supported additional platform services, while Datastream and BigQuery helped simplify data movement and analysis.

Outcomes

- Writer.com **handles over a trillion API calls monthly** with minimal database overhead
- The migration to AlloyDB for PostgreSQL **reduced operational complexity** and freed up technical resources
- Platform performance scaled to **support 90,000 words generated per second**

[Read the full story](#)

“Google Cloud databases allow us to manage a highly sophisticated and complex platform securely. Our ability to be highly effective with our technical talent and scale our output efficiently is an incredible competitive advantage.”

J.R. ROBINSON

Senior Director Infrastructure, Platform and Compliance, Writer.com

**Industry**

Startups and digital natives

Country

United States

Google Cloud database[Spanner](#)

Prefab gains 100x more storage with no maintenance downtime

Challenge

Prefab helps developers ship apps faster with feature flags, dynamic logging, and secrets management. The company set out to find a database that could deliver the versatility and incredible scale needed for its feature flag service.

Solution

Prefab adopted Spanner's PostgreSQL interface to support high-throughput APIs and large-scale telemetry tracking with strong consistency, availability, and horizontal scaling. Spanner handled unpredictable data growth with no downtime, making it easy to increase performance by adjusting processing units. With built-in multi-region replication and familiar PostgreSQL syntax, Spanner gave Prefab the operational reliability of a NoSQL system with the usability of a relational database. Prefab used Key Visualizer to troubleshoot hotspots and tune performance over time.

Outcomes

- **Infrastructure costs dropped** to one-third of comparable PostgreSQL setups
- The system **scaled to 37 million rows** with no downtime or re-architecting
- Spanner delivered **99.99% uptime and multi-zone resilience** for critical workloads

[Read the full story](#)

“Spanner is easy to set up, easy to use, and – surprisingly – less expensive than other databases we’ve tried for workloads that need the option to scale. We’re already impressed by the performance to date, and we’re nowhere close to its limits yet.”

JEFF DWYER

Founder & CEO, Prefab

**Industry**

Startups and digital natives

Country

Spain

Google Cloud databases[AlloyDB for PostgreSQL](#),[Firestore](#)**Additional Google****Cloud products**App Engine, BigQuery, Cloud
Armor, Cloud Load Balancing

Loyal Guru slashes query latency by up to 50% with AlloyDB for PostgreSQL

Challenge

Loyal Guru, a loyalty and offer personalization platform, needed to process and analyze large volumes of customer data in real time to deliver tailored promotions. However, database bottlenecks during high-traffic events slowed query performance and led to delays in customer engagement.

Solution

Loyal Guru migrated its transactional workloads to AlloyDB for PostgreSQL, which granted low-latency access to critical data across customer profiles, loyalty transactions, and offer redemptions. In addition, Firestore supported high-concurrency reads for static data, while BigQuery handled large-scale behavioral analytics.

Outcomes

- AlloyDB for PostgreSQL **reduced query latency by up to 50%** and **optimized storage by 35%**
- AlloyDB for PostgreSQL **automatically adjusts storage capacity to meet demands** during traffic spikes without manual intervention
- Database response times for **API requests have improved by 60%**

[Read the full story](#)

“AlloyDB for PostgreSQL, along with other Google Cloud products like Firestore and BigQuery, provides the agility and performance needed to continually enhance our platform’s capabilities and help us anticipate emerging trends rather than merely reacting.”

JESÚS ANTONIO CANALES DIEZ

Platform Tech Lead, Loyal Guru

KOCHAVA★

Industry

Startups and digital natives

Country

United States

Google Cloud database

[Spanner](#)

Additional Google

Cloud products

BigQuery, Google Kubernetes Engine, Pub/Sub

Kochava powers real-time mobile analytics at scale with Spanner

Challenge

Kochava, a leading mobile analytics and attribution platform, needed a way to process millions of ad signals per minute per customer without bottlenecks or inconsistency. Its mix of on-prem databases created fragmented experiences and couldn't scale with demand.

Solution

Kochava consolidated real-time data storage on Spanner to eliminate database fragmentation and deliver a consistent, high-performance experience. Each customer's database is provisioned automatically in Spanner, scaling horizontally with autoscaling and infrastructure-as-code. Developers benefit from familiar SQL and built-in features like TTL, while BigQuery supports reporting needs.

Outcomes

- Spanner's autoscaler has enabled Kochava to **optimize costs and reduce consumption by 30%**
- The Kochava team can **focus on building applications** rather than addressing limitations of its legacy database
- Spanner supports time to live, which **eliminates running background programs** to see if data can be deleted, reducing table size

[Read the full story](#)

“We tried various databases, and Spanner was the easiest to implement, had the best performance, and offered the best managed service.”

NICK OTTER

Lead Software Engineer, Kochava

**Industry**

Startups and digital natives

Country

United States

Google Cloud database[Bigtable](#)

Bigtable enables Airship to write 1 million row operations per second

Challenge

Airship, which provides a platform for building customer experiences, was struggling to manage its HBase instance efficiently because of the operational challenges it presented. The team wanted to lower their operational burden and started evaluating fully managed database services.

Solution

Airship replaced HBase with Bigtable, Google Cloud's fully managed NoSQL database. Bigtable's HBase client wrappers enabled a smooth migration, allowing Airship to dual-write during the transition without data conflicts. Now, instead of maintaining clusters, Airship simply scales performance with a few clicks.

Outcomes

- One of Airship's clusters is reading **700,000 rows per second**
- Airship has **freed up developers to focus on CI/CD and tooling** to create high-value features that customers love
- Airship can **easily write a generic job** to load a SequenceFile of rows snapshotted from HBase into Bigtable

[Read the full story](#)

“As Airship grows, we plan to scale all Google Cloud services to keep our system both performant and cost effective. We consider Google Cloud a key partner when it comes to our technology that helps brands master mobile app experiences.”

NEIL GARIEPY

VP Engineering & Security, Airship

**Industry**

Startups and digital natives

Country

United States

Google Cloud databases[Bigtable](#), [Memorystore](#)**Additional Google****Cloud products**

BigQuery, Google Kubernetes Engine, Vertex AI

Moloco handles more than 5 million ad bid requests per second with Bigtable

Challenge

Using strategic mobile ad placement, Moloco develops machine learning solutions to help app advertisers acquire, re-engage, and retain high-value mobile app users. Moloco needed a low-latency, scalable solution to respond to millions of bid requests in less than 150 milliseconds each.

Solution

Moloco built its real-time bidding platform on Bigtable, using it as persistent storage for campaign and bidding history. Memorystore provides fast in-memory access to frequently queried data, complementing Bigtable's low-latency reads. The system scales to meet daily traffic surges and enables access to billions of sparse data records. Bid requests pass through GKE-hosted services, triggering inferences from TensorFlow models powered by data from Bigtable and contextual features.

Outcomes

- Bigtable scaled bid handling capacity from 550,000 to **over 5 million requests per second**
- Moloco can autoscale Bigtable instances without impacting reliability, **improving efficiency and reducing overhead**
- Low-latency access and long-term, persistent **storage for billions of records** in Bigtable are used for machine learning

[Read the full story](#)

“With the help of Google Cloud services, Moloco can now process hundreds of billions of bid requests every day [...] Bigtable offers low latency and high scalability to access our history data, and we're able to transform our data and our customers' data into valuable insights.”

CHANG KIM

VP of Engineering, Moloco

**Industry**

Startups and digital natives

Country

United States

Google Cloud database[Memorystore for Redis Cluster](#)**Additional Google Cloud products**

Google Cloud VPC

Statsig supports up to 7.5 million queries per second with Memorystore for Redis Cluster

Challenge

Statsig helps companies ship, test, and manage software and application features with confidence. Facing bottlenecks and connectivity issues, the Statsig team realized they needed a performant, reliable, scalable, and fully managed Redis service.

Solution

Statsig migrated to Memorystore for Redis Cluster to support stateless applications with dynamic workloads. The team uses Memorystore to power core services including real-time health checks, feature flag delivery, and streaming event deduplication. Memorystore's high availability and built-in monitoring tools eliminated recurring Redis connection issues and simplified operations.

Outcomes

- **Increased efficiency led to 70% lower cost** of Redis compared to the costs of running the same workloads on previous cloud provider
- Statsig can **easily handle an average of 1.5 million QPS and up to 7.5 million QPS at peak**, keeping pace with customers as they scale
- Ability to **scale in or out with zero downtime**, enabling support for higher QPS and a range of use cases that will drive future growth

[Read the full story](#)

“Memorystore for Redis Cluster has allowed us to accomplish our business goals without compromising on cost or predictability.... It has become an invaluable asset, delivering robust scalability and versatility for our operations.”

JASON WANG

Software Engineer, Statsig

Healthcare and life sciences

Turn data into scientific breakthroughs and improved patient outcomes with fast, reliable, and compliant data solutions.



**Industry**

Healthcare and life sciences

Country

Germany

Google Cloud database[AlloyDB for PostgreSQL](#)**Additional Google****Cloud products**

Datastream, BigQuery, Google Kubernetes Engine, Pub/Sub, Vertex AI

Bayer increases throughput by 5x and cuts response times in half with AlloyDB for PostgreSQL

Challenge

Bayer uses the power of science to shape the future of farming. Its Global Data Assets team created Field Answers – a modern data solution that stores and analyzes vast amounts of observational data. Extensive load testing revealed that its current database was inadequate for onboarding new market segments, leading to potential bottlenecks and replication lag.

Solution

Bayer migrated its data operations to AlloyDB for PostgreSQL, eliminating bottlenecks by centralizing replication through a single source of truth. AlloyDB's compatibility with PostgreSQL allowed Bayer to transition without modifying applications, significantly cutting down the migration timeline and ensuring stability during the crucial planting season.

Outcomes

- Migrating to AlloyDB for PostgreSQL **reduced response times by over 50%** on average
- Bayer **increased throughput by 5x** compared to the previous PostgreSQL setup
- AlloyDB for PostgreSQL's single source of truth for all nodes **significantly reduced the impact of scaling read traffic**

[Read the full story](#)

“By migrating to AlloyDB for PostgreSQL, we’ve ensured that our business growth won’t be hindered by database limitations, allowing us to focus on innovation.... By leveraging the power of AlloyDB for PostgreSQL and the Google Cloud ecosystem, we’re not only enhancing our own operational capabilities but also contributing to the future of farming.”

AARON JOYCE

Engineering Lead, Bayer Crop Science



Industry

Healthcare and life sciences

Country

United States

Google Cloud database

[Cloud SQL for PostgreSQL](#)

[Enterprise Plus edition](#)

Additional Google

Cloud products

Google Compute Engine

Intelligencia AI doubles deployment speed with Cloud SQL

Challenge

Powered by AI, Intelligencia AI provides a suite of solutions to help businesses in the pharmaceutical space make more informed, data-driven decisions to de-risk clinical development. With skyrocketing data volumes, commitment to quality, and complex processing requirements, the company realized it needed a robust, scalable, and reliable database.

Solution

With its seamless integration, high availability, and automatic backups, Cloud SQL for PostgreSQL Enterprise Plus edition provided Intelligencia AI with a powerful tool to efficiently store and manage data in real time.

Outcomes

- Intelligencia AI **reduced deployment time by 50%** through automated deployment of feature environments
- Effortless resource scaling and networking configuration changes **increased operational agility and efficiency**
- Restoring backups to multiple instances allows for **continuous deployment of data generated by ETL pipelines to production**

[Read the full story](#)

“Cloud SQL for PostgreSQL Enterprise Plus edition provides the reliability and scalability needed for our AI-driven probability of technical regulatory success (PTRS) predictions, ensuring accurate and transparent results for our customers.”

ALBERTO BENROUBI

Principal Data Engineer, Intelligencia AI

**Industry**

Healthcare and life sciences

Country

United States

Google Cloud database[AlloyDB for PostgreSQL](#)**Additional Google Cloud products**

Cloud GPUs, Compute Engine, Google Cloud Storage, Vertex AI

NeuroPace uses AlloyDB AI to enhance seizure identification

Challenge

NeuroPace, a medical device company dedicated to transforming epilepsy treatment, collected over 15 million intracranial electroencephalogram (iEEG) records. Its on-premises infrastructure restricted the speed of machine learning (ML) model optimization, hindering fast identification of seizure patterns crucial for personalized patient care.

Solution

NeuroPace migrated its ML workloads to Google Cloud, significantly accelerating model training and optimization. Using AlloyDB AI's built-in vector search capabilities, the company quickly identifies similar iEEG patterns across millions of records. Integration with Vertex AI and Compute Engine streamlined model training, offering greater efficiency, scalability, and faster insights.

Outcomes

- NeuroPace **accelerates the speed of ML training with AlloyDB AI**, searching through millions of records in milliseconds—not hours
- AlloyDB AI allows NeuroPace to store data embeddings in vector form, **simplifying and expediting similarity searches**
- Vertex AI and Google Cloud GPUs offer NeuroPace a **better price-performance ratio compared to on-premises**

[Read the full story](#)

“Google Cloud’s technologies have significantly improved and accelerated NeuroPace’s ML training capabilities. Searching through more than a million iEEG records to identify similar ones, a task that previously took minutes to hours, can now be completed in milliseconds using Google’s AlloyDB AI.”

SHARANYA DESAI, Ph.D.

Technical Fellow and Director of AI, NeuroPace, Inc.

Games

Real-time action needs real-time data. Keep players engaged and your game running strong.



**Industry**

Games

Country

United States

Google Cloud database[Cloud SQL for MySQL](#)[Enterprise Plus edition](#)**Additional Google****Cloud products**

Database Migration Service

Chess.com boosts performance and reduces latency with Cloud SQL

Challenge

Chess.com is the premium platform for online chess and one of the largest online games in the world. After experiencing a major surge in demand, it needed to adopt a scalable, global architecture to meet high performance and SLA requirements.

Solution

An upgrade to the scalable architecture of Cloud SQL for MySQL Enterprise Plus edition resolved Chess.com's server capacity challenges, reduced operational burdens, and enhanced its user experience. By deploying databases in several Google Cloud regions, it now seamlessly connects millions of chess players across the globe.

Outcomes

- The **p99 latency response was reduced by 71.4%** – going from 14ms to an impressive 4ms
- A reduced server footprint significantly **cut costs and improved maintenance by 90%**
- New optimizations **reduced task times from 5–10 minutes to a few seconds**

[Read the full story](#)

“We’re thrilled about the seamless capabilities that Google Cloud databases offer. The ease with which we can manage and create new database instances aligns perfectly with our agile development approach.”

GRZEGORZ DLUGOLECKI

Principal Cloud and Kubernetes Engineer, Chess.com

**Industry**

Games

Country

Japan

Google Cloud database[Spanner](#)**Additional Google****Cloud products**BigQuery, Cloud Storage,
Google Kubernetes Engine

COLOPL optimizes player and developer experience with Spanner

Challenge

Mobile gaming application developer COLOPL wanted to reduce costs and increase the scalability and availability of its mobile gaming applications, while freeing the business to focus more on game development.

Solution

With Spanner, COLOPL eliminated the scalability constraints that occurred when using horizontally and vertically partitioned databases for large-scale services. Spanner has also enabled COLOPL to seamlessly and cost-effectively change the number of nodes in its service daily, depending on the needs of the business.

Outcomes

- COLOPL **eliminated scalability constraints** by merging all user data into a single database
- **Developers were able to focus more on game logic** instead of on managing the complexities of multiple databases
- **Scaling nodes daily** reduces the cost of adding excess server capacity and ensuring availability

[Read the full story](#)

“Depending on the number of databases the organization needed to scale, making the adjustments needed to respond to big load changes could take up to 5.5 man-days. Spanner has helped the organization reduce database costs by up to 25% and operational costs by up to 80%.”

COLOPL TEAM

Manufacturing and supply chain

The future of production and delivery is connected, data-driven, and built to scale.



**Industry**

Manufacturing and supply chain

Country

United States

Google Cloud databases

[Cloud SQL](#), [AlloyDB for PostgreSQL](#), [Spanner](#), [Memorystore](#), [Firestore](#)

Additional Google Cloud products

Cloud Run, MongoDB Atlas on Google Cloud

Ford reduces operational workload with Cloud SQL

Challenge

Ford Motor Company, one of the most recognizable auto brands in the world, needed to modernize its mixed fleet of on-premises and cloud databases. Routine management tasks like provisioning, patching, and scaling were time-consuming and unpredictable, requiring a global team of database administrators.

Solution

Ford migrated its databases to fully managed Google Cloud services, using an opinionated stack of migration tools to streamline the move and support future application development. With serverless tools like Cloud Run and Cloud SQL, Ford improved database processing while eliminating manual backup management. Cloud SQL helped Ford meet performance and protection requirements with minimal downtime and reduced operational burden.

Outcomes

- Some workloads experienced a **30% performance boost** post-migration
- Ford **reduced the time spent on database-related operational tasks** across its global teams
- **Zero backup failures** across migrated Cloud SQL databases improved data reliability

[Read the full story](#)

“At the database level, our goal is to enable always-on products with minimum downtime. By migrating to fully managed Google Cloud databases like Cloud SQL, we significantly reduced our management overhead. We’ve already seen a large drop in database-related operational tasks.”

NIMISHA SHAH

Director of Database Services, Enterprise Platform
Engineering and Operations at Ford Motor Co.

**Industry**

Manufacturing and supply chain

Country

United States

Google Cloud database

[Bigtable](#)

Additional Google**Cloud products**

BigQuery, Compute Engine, Dataflow, Pub/Sub

Ford Pro delivers real-time vehicle insights with Bigtable

Challenge

Ford Pro Intelligence is a platform for managing commercial fleets, serving customers from small businesses to large municipal operations. As vehicle telemetry grew in complexity and volume, Ford needed a database that could handle high-throughput, low-latency access for both real-time and historical data across a dynamic schema.

Solution

Ford Pro built its telemetry platform on Bigtable, a NoSQL database built for scale and time-series data. Incoming data from vehicle sensors is streamed using Pub/Sub and Dataflow, then written to Bigtable for operational and historical queries. Bigtable's built-in support for sparse data, schema flexibility, and lifecycle management helped reduce cost and overhead. Real-time data powers customer-facing APIs and dashboards for tracking fleet health, vehicle activity, and location.

Outcomes

- Bigtable supports **real-time APIs with single-digit millisecond latency** across thousands of fleets
- Time-series architecture enables **predictive maintenance and live vehicle health alerts**
- Fully managed services **reduced infrastructure burden and improved compliance at scale**

[Read the full story](#)

“We knew that we needed an operational data store that could support low-latency access for both real-time and historical data with a flexible schema. In the end, the choice was obvious, and we decided to use Bigtable as our central vehicle telemetry data repository.”

GAVARRAJU NANDURI

Head of Data Engineering, Ford Pro

Renault Group

Industry

Manufacturing and supply chain

Country

France

Google Cloud database

[Cloud SQL](#)

Additional Google

Cloud products

BigQuery, Dataflow

The Renault Group completes ambitious 70-app migration on Cloud SQL

Challenge

Renault, known for its iconic French cars, needed to modernize 70 custom-built Quality and Customer Satisfaction applications spanning decades of technology. To achieve this ambitious goal, the team wanted a more agile, cost-effective alternative to on-premises Oracle databases.

Solution

Renault migrated all its applications to Google Cloud in just two years, using Cloud SQL for PostgreSQL as the foundation. The team started with low-complexity apps, tracked SQL patterns in an internal wiki, and optimized legacy code and queries during replatforming. Built-in monitoring tools helped pinpoint performance bottlenecks and fine-tune infrastructure. Cloud SQL gives teams the flexibility to increase compute power or reduce infrastructure as needed – no hardware provisioning required.

Outcomes

- Renault completed the [migration of 70 in-house applications from Oracle to Cloud SQL in just two years](#)
- Batch processing time was reduced by 66%, [cutting nine-hour jobs down to three hours](#)
- The migration has [lowered costs down to one dollar per user per year](#) because there is no overprovisioning

[Read the full story](#)

“The approach we developed was very successful – where database migration was initially seen as insurmountable, the entire migration project was completed in two years.”

CYRIL PICCHIOTTINO

VP, Quality & Customer Satisfaction IS, The Renault Group



Industry
Manufacturing and supply chain

Country
India

Google Cloud databases
[Spanner](#), [Memorystore](#)

Additional Google Cloud products
Cloud Run functions, Google Kubernetes Engine, Pub/Sub

Mahindra builds a first-of-its-kind digital SUV order system with Spanner

Challenge

As one of India's largest automotive companies, Mahindra wanted to give customers a new way to book vehicles during the launch of its latest SUV. Traditional, dealer-led launches couldn't support the scale and speed Mahindra anticipated from its digital-first buyers.

Solution

Mahindra built a cloud-first booking platform designed to support massive bursts of traffic with high concurrency and low latency. Spanner handled high-throughput transactions and real-time sales visibility, while Google Kubernetes Engine and Memorystore scaled automatically during the spike in activity. The solution used Pub/Sub and Cloud Run functions to deliver booking confirmations in real time and handle extensive load testing to prepare for 250,000 concurrent users.

Outcomes

- Mahindra **processed 25,000 booking requests in the first minute** and 100,000 in the first 30 minutes
- At peak, the system supported **60,000 concurrent users without performance issues**
- The digital booking window helped Mahindra secure **\$2.3 billion in SUV orders on launch day**

[Read the full story](#)

“Using Spanner not only provided us the scale needed to store the car bookings rapidly, but also allowed the admin teams to see real-time drill-down pivots of sales performance across vehicle models, towns, and dealerships, without the need for an additional analytical processing layer.”

BHUWAN LODHA
Senior VP and Chief Digital Officer, Automotive, Mahindra Group



Industry

Manufacturing and supply chain

Country

United States

Google Cloud database

[Cloud SQL](#)

Manhattan Associates boosts uptime and reduces costs with Cloud SQL

Challenge

Manhattan Associates powers mission-critical supply chain solutions for global retailers, requiring always-on performance and minimal operational overhead. Its legacy database systems were expensive, inflexible, and difficult to manage across cloud environments.

Solution

To support its Manhattan Active® Platform, the company migrated to Cloud SQL for MySQL, a fully managed database service that simplifies administration and delivers high availability. Cloud SQL reduced replication lag, accelerated replica creation, and enabled the team to resize databases in minutes. The entire migration from the previous cloud vendor was completed with less than four hours of downtime.

Outcomes

- **Unplanned downtime dropped by 83%** compared to the previous solution
- The team now **runs hundreds of Cloud SQL instances** with only a few DBAs
- Replica creation for TB-scale databases now takes **under 30 minutes instead of several days**

[Read the full story](#)

“Cloud SQL provides highly scalable, available, and reliable database capabilities within Manhattan Active Platform, which helps us provide significantly better outcomes for our clients and better experiences for their customers.”

SANJEEV SIOTIA

SVP and CTO, Manhattan Associates



Industry

Manufacturing and supply chain

Country

United States

Google Cloud database

[Bigtable](#)

Additional Google

Cloud products

Compute Engine, Google Kubernetes Engine, Pub/Sub

Oden Technologies improves smart factory efficiency with Bigtable

Challenge

Oden Technologies helps manufacturers use real-time data to optimize factory operations. To cut costs and simplify infrastructure, Oden looked to replace its existing cloud provider while improving performance and scalability.

Solution

Oden migrated its platform to Google Cloud and leveraged Bigtable to process and retrieve tens of millions of factory metrics daily. The team built dashboards that pull real-time insights from Bigtable, helping customers understand performance trends, troubleshoot faster, and optimize production lines. The move to Bigtable also allowed Oden to streamline infrastructure and reduce overhead.

Outcomes

- Reduced VM instance count from 80 to 45, **cutting infrastructure complexity**
- **Lowered costs of storage and analytics by 30%**
- Delivered **real-time visibility into factory performance**, helping manufacturers reduce downtime and waste

[Read the full story](#)

“With Google Cloud, we are helping our customers to be data-driven, which wasn’t possible before. They now understand that data is their most important asset. That allows them to be more innovative and continually improve their production processes.”

WILLEM SUNDBLAD

CEO and Founder, Oden Technologies



Industry

Manufacturing and supply chain

Country

India

Google Cloud database

[Bigtable](#)

Additional Google

Cloud products

BigQuery, Cloud Run functions, Dataflow, Pub/Sub

Bharat Light & Power improves asset health with Bigtable

Challenge

Clean energy company Bharat Light & Power (BLP) needed a way to improve wind turbine reliability and industrial equipment uptime across sectors such as manufacturing, ports, and rail. A single failed component could halt operations and lead to major losses.

Solution

BLP built its Orion platform on Google Cloud, using Bigtable to store and retrieve high-volume telemetry data from thousands of turbines and sensors. The NoSQL database supports real-time monitoring and failure prediction by handling continuous input from diverse industrial systems. Combined with BigQuery for analysis and Cloud Run functions for automation, the architecture helps BLP customers detect issues early and respond quickly to keep operations running.

Outcomes

- Predictive maintenance capabilities were **deployed across 2,000 wind turbines in multiple countries**
- Manufacturing customers **increased productivity by at least 5% and reduced costs by 10%**
- The system **processes over 8 million turbine data points and 200,000 sensor data points each day**

[Read the full story](#)

“Our technology strategy is to connect the most advanced industrial IoT hardware to software and the cloud – using the most powerful AI cloud available. That’s why we chose to bring our AI algorithms to Google Cloud.”

TEJPREET SINGH CHOPRA

Chief Executive Officer, Bharat Light & Power

**Industry**

Manufacturing and supply chain

Country

India

Google Cloud database

[Bigtable](#)

Additional Google**Cloud products**

BigQuery, Dataflow, Google Kubernetes Engine, Identity and Access Management (IAM), Pub/Sub

[Read the full story](#)

STL unlocks real-time factory insights with Bigtable and BigQuery

Challenge

Optical fiber company STL needed to orchestrate complex global operations – from manufacturing to network deployment – while maintaining high margins in a competitive market. Data across teams and systems was fragmented, limiting the ability to analyze performance and optimize processes.

Solution

STL built a data lake on Google Cloud with BigQuery and Bigtable at its core. BigQuery supports scalable analytics across global operations, providing predictive insights that inform manufacturing processes and complex project management. Bigtable stores IoT data from factory equipment, allowing rapid access and analysis within milliseconds. This assists engineers in identifying issues affecting yield and quality.

Outcomes

- BigQuery delivers **real-time insights in under 5 seconds**, accelerating decision-making across global teams
- Predictive models built on BigQuery **improved on-time delivery and protected margins** for complex deployments
- Bigtable **processes IoT data from factory equipment in milliseconds**, helping engineers improve manufacturing yield


“We managed to build ML models in just months instead of years, to access granular data from business processes. This enabled us to augment human intelligence by providing the process engineering and R&D teams with insights that were not previously known.”

MANUJ DESAI

Head of IT Transformation, STL

Security

Scale your defenses with a database that handles massive throughput, spikes in demand, and global operations without breaking a sweat.





Industry
Security

Country
United States

Google Cloud database
[Bigtable](#)

Additional Google Cloud products
Dataflow, Pub/Sub

Palo Alto Networks reduces cost and latency by migrating to Bigtable

Challenge

Palo Alto Networks needed a database that could support the scale and uptime requirements of Advanced WildFire, its cloud-based malware analysis and prevention solution. One key component, the Global Verdict Service, delivers real-time threat assessments based on WildFire's analysis. As usage grew, Cassandra's latency, operational complexity, and scaling limitations couldn't keep up with demand.

Solution

To improve performance and simplify operations, Palo Alto Networks migrated the Global Verdict Service to Bigtable. The team used a phased approach with dry-run and final migrations, dual writes, and data integrity checks to maintain consistency throughout. Bigtable's fully managed architecture, global scalability, and sub-millisecond performance eliminated the need for a complex mesh replication setup and reduced operational overhead.

Outcomes

- Bigtable **reduced read/write latency by 5x** to significantly improve application responsiveness
- Moving to a managed service model **cut database-related costs by 50%**
- Availability jumped from 99.95% to an impressive 99.999%, **ensuring near-constant uptime and minimizing service disruptions**

[Read the full story](#)

“Palo Alto Networks’ migration to Bigtable was transformative. If your organization is grappling with database challenges like performance bottlenecks, scalability limitations, or operational complexity, consider Bigtable.”

RAVI PARUCHURI

Sr. Principal Engineer, Palo Alto Networks



Industry
Security

Country
United States

Google Cloud database
[Cloud SQL for PostgreSQL](#)

**Additional Google
Cloud products**
BigQuery

NetRise accelerates vulnerability detection with Cloud SQL

Challenge

Cybersecurity company NetRise needed a scalable, high-performance backend to support its Trace platform, which detects vulnerabilities in complex software supply chains across embedded and cyber-physical systems.

Solution

NetRise adopted Cloud SQL for PostgreSQL with pgvector to power Trace, combining relational storage with pgvector for fast, natural language-based semantic search. This supported storing vector embeddings to detect issues like hardcoded credentials across complex, nested assets. The team later added BigQuery to accelerate large-scale analytics, replacing Elasticsearch and streamlining infrastructure management to support a machine learning-driven approach to cybersecurity.

Outcomes

- Cloud SQL for PostgreSQL with pgvector **reduced server resource needs by 50% and cut query response times by 60%**
- NetRise now delivers **10x faster threat research and detection** for customers and internal teams
- Switching to BigQuery accelerated data processing 30x, **cutting a 24-hour workload down to 47 minutes**

[Read the full story](#)

“Turning to Google Cloud’s managed services was a game-changer. Cloud SQL has been pivotal in scaling our architecture and optimizing queries, significantly reducing the time and resources required for complex data analysis.”

MICHAEL SCOTT
CTO and Co-Founder, NetRise



Industry
Security

Country
Denmark

Google Cloud database
[AlloyDB for PostgreSQL](#)

Tricent Security Group realizes up to 25% savings with AlloyDB for PostgreSQL

Challenge

Tricent Security Group, a provider of file-sharing security tools, needed a scalable solution to support both online transaction processing workloads – like real-time file permission updates – and complex online transaction processing queries for its analytics platform. Its PostgreSQL setup on VMs struggled to keep up, limiting performance and growth.

Solution

Tricent chose to migrate its database operations to AlloyDB for PostgreSQL to enable efficient handling of both online transaction processing (OLTP) and online analytics processing (OLAP) workloads—without needing to rewrite database code. AlloyDB's columnar engine significantly accelerated analytical queries, while its separation of compute and storage provided flexibility to quickly scale resources up or down based on real-time demands.

Outcomes

- System now handles **250 million transactions per day without latency**
- **Month-over-month cost savings of 10–25%** due to improved efficiency
- **Freed up engineering teams** to focus on product innovation instead of infrastructure

[Read the full story](#)

“AlloyDB for PostgreSQL’s unique architecture separates compute and storage, giving us the flexibility to adapt quickly to changing demands. This kind of elasticity is invaluable when dealing with unpredictable workloads and rapid growth.”

JAKOB TOLSTRUP BECH

Chief Product and Technology Officer, Tricent Security Group A/S



Industry
Security

Country
United States

Google Cloud database
[Bigtable](#)

Additional Google Cloud products
Dataflow

Stairwell powers real-time threat detection and analysis with Bigtable

Challenge

Cybersecurity company Stairwell needed to store and analyze all executable files across customer environments – indefinitely. Its original PostgreSQL database couldn't support the volume, performance, or scaling requirements of its growing workload.

Solution

Stairwell migrated its key-value data storage to Bigtable to support real-time lookups, fast scans, and large-scale batch processing. Bigtable underpinned both the threat analysis platform and the machine learning-based malware detection system by delivering low-latency access and high-throughput processing without downtime.

Outcomes

- A single table housing hundreds of billions of data points maintained an **average read latency of just 1.9ms**
- Automated scaling and node count adjustments for query loads **ensured consistent performance**
- In one instance, Bigtable effortlessly **served more than 22 million rows per second** during an intense Dataflow job

[Read the full story](#)

“Bigtable isn't just a gigantic storage vault for us; it's a high-performance analytics engine capable of executing both batch and streaming queries on a grand scale. It provides a game-changing level of data processing capability.”

MIKE WIACEK

Chief Executive Officer and Founder, Stairwell



Industry
Security

Country
United Kingdom

Google Cloud database
[Bigtable](#)

Additional Google Cloud products
BigQuery, Google Kubernetes Engine, Elastic on Google Cloud

Ravelin detects retail fraud with low latency using Bigtable

Challenge

As a fraud detection platform for online retailers, Ravelin needed a highly scalable database that could support low-latency access to customer and transaction history during checkout. Its existing system couldn't keep up with growing traffic and the demands of larger clients.

Solution

Ravelin migrated from AWS to Bigtable, adopting a key-value structure for fast, secure lookups. Every fraud decision relies on real-time access to customer history in Bigtable, which quickly ingests and processes high volumes of data. The team also used BigQuery for analytics and Elastic on Google Cloud for dashboard search. A mirrored traffic approach ensured a seamless migration with no downtime.

Outcomes

- Bigtable powers **22,000+ requests per second** in production with low latency
- Scaling to 167,000 requests per second during load testing required **only a single config change**
- The **highest level of security and data protection** keeps customers' data safe

[Read the full story](#)

“When a client’s customer places an order, we need to process their full history and as much data as possible about that customer in order to detect fraud, all while keeping their data secure. Bigtable excels at accessing and processing that data in a short time window.”

JONO MACDOUGALL

Software Engineer and Co-Founder, Ravelin

Media and entertainment

Spark and hold audience attention with personalized content, real-time insights, and always-on data.



**Industry**

Media and entertainment

Country

United States

Google Cloud database[Bigtable](#)

YouTube supports billions of creators and views with Bigtable

Challenge

As one of the world's largest streaming platforms, YouTube needed a way to store and serve metadata on billions of videos, channels, and playlists to power creator dashboards, payments, and analytics. The system had to support real-time ingestion, historical tracking, and querying at massive scale.

Solution

YouTube built a metadata warehouse on Bigtable to drive reporting and analytics across its platform. Bigtable's flexible schema let teams land raw data quickly, then evolve models as needs changed. The architecture ingests raw data from canonical sources, transforms it, and serves it to downstream clients. To keep data fresh, YouTube built pipelines that detect changes to source data and automatically pull in the latest information for reporting.

Outcomes

- Bigtable **supports analytics for billions of videos and creators** across YouTube
- Data powering creator dashboards is **updated automatically as source systems change**
- Bigtable provides **high performance at low cost** across mixed workloads

[Read the full story](#)

“Bigtable has transformed the way we handle data at YouTube. It's not just a database, it's an enabler of innovation, agility, and customer-centricity. It's the cornerstone of our evolving global infrastructure.”

BIN LIU

Software Engineer, YouTube

Forbes

Industry

Media and entertainment

Country

United States

Google Cloud database

[Firestore](#)

Additional Google

Cloud products

BigQuery

Forbes publishes high-performing content using Firestore

Challenge

Global media company Forbes needed a modern way to process site statistics for contributors. Its legacy on-prem system was complex, costly, and offered no control over contributor access.

Solution

Forbes replaced its custom-built logging pipeline and MySQL setup with Firestore to power a real-time statistics system that's fully cloud-native and maintenance-free. Firestore integrated directly with Google Analytics and BigQuery, giving contributors fast, granular access to performance data. The new system reduced infrastructure from 40+ application servers to just three and added new capabilities like automated SEO recommendations.

Outcomes

- **Site metrics are now refreshed every 15 minutes** instead of once daily
- Writers receive **real-time performance data**, down to the minute
- Infrastructure consolidation **cut application server count by over 90%**

[Read the full story](#)

“Once we implemented our new statistics processing system, we were able to update our contributors’ site metrics much faster.... By providing this granular level of data to our contributors, we are helping them better optimize their content and deliver the best possible pieces to their readers.”

BENJAMIN HARRIGAN

Software Architect, Forbes

GRUPO GLOBO

Industry

Media and entertainment

Country

Brazil

Google Cloud database

[Bigtable](#)

Additional Google

Cloud products

BigQuery, Dataflow, Pub/Sub

Grupo Globo cuts streaming infrastructure costs by 60% with Bigtable

Challenge

As Latin America's largest media group, Grupo Globo runs the Globoplay streaming service for live and on-demand content. Its "Continue Watching" API, which handles high-volume write traffic, relied on Cassandra – but scaling to meet demand was manual, expensive, and time-consuming.

Solution

Globo migrated its Cassandra-based system to Bigtable to take advantage of its managed operations, autoscaling, and high availability. The team gradually transitioned the write path and built a batch pipeline in Dataflow to migrate legacy data. With 100% of traffic redirected to Bigtable, Globo retired Cassandra and rolled out Pub/Sub to further modernize background processing.

Outcomes

- Bigtable **reduced infrastructure costs by approximately 60%**
- Eliminating manual scaling **cut operational overhead and improved reliability**
- Streaming service now delivers read-your-writes consistency and **high performance under heavy traffic**

[Read the full story](#)

“Migrating to Bigtable decreased our maintenance needs and gave us guaranteed database scalability. We’re excited to continue this partnership in simplifying database management in order to meet our business’ ever-evolving demands.”

MICHEL HENRIQUE AQUINO SANTOS

Software Engineer, Grupo Globo

FANCODE**Industry**

Media and entertainment

Country

India

Google Cloud database[Memorystore](#)

FanCode doubles live sports streams with Memorystore for Redis Cluster

Challenge

FanCode, India's leading sports streaming platform, needed to deliver personalized, real-time content to millions of fans across devices. As traffic and content volume grew, its legacy Redis setup introduced latency and scaling issues that hurt performance during live events.

Solution

FanCode migrated from self-hosted Redis on AWS to Memorystore for Redis Cluster in order to align with its Google Cloud infrastructure. Memorystore now powers the caching layer across FanCode's microservices, supporting low-latency data delivery for features like live scores. Memorystore's fully managed and scalable architecture lets FanCode expand clusters on demand – without major reconfiguration.

Outcomes

- FanCode now **supports over 15,000 live events per year** – more than double the previous year
- The platform **streams billions of minutes of live content annually with sub-second latency**
- A fully managed caching layer **reduced operational overhead** for FanCode's small team

[Read the full story](#)

“With Memorystore for Redis Cluster, we built a fully integrated, scalable infrastructure that powers real-time fan experiences during high-traffic sports events.”

AMIT MIRCHANDANI

CTO, FanCode

**Industry**

Media and entertainment

Country

India

Google Cloud database[Spanner](#)

Glance reduces latency with Spanner

Challenge

Powered by generative AI, the Glance app delivers live and tailored content to more than 200 million users. Glance was experiencing cumbersome operational and data storage issues with its legacy database and required a solution that was more agile, resilient, and reliable.

Solution

Glance migrated to Spanner because of its ability to meet required transaction scale, maintain low latency, and simplify relational data management. Spanner's relational semantics and horizontal scalability allowed Glance to efficiently handle complex JOIN operations and schema updates without downtime. The fully managed database service also simplified schema evolution for faster and more frequent feature releases.

Outcomes

- **Day-to-day database operations have improved**, providing multiple monitoring panels and insights for queries, transactions, and locks
- Many schema updates now run as background tasks with no impact to the live database, resulting in **zero downtime**
- Latency significantly improved, with **less than 13ms for p99 latencies in read and write**

[Read the full story](#)

“Spanner improved our day-to-day database operations. After implementation, we achieved the required scale of queries per second that we needed, and reached client-side latencies below 20ms for both read and write.”

HARDIK TALUJA

Software Development Engineer III, Glance

**Industry**

Media and entertainment

Country

United States

Google Cloud databases[Bigtable](#), [Memorystore for Memcached](#)**Additional Google****Cloud products**

BigQuery, Google Kubernetes Engine, Pub/Sub

OpenX serves 150 billion daily ad requests with Bigtable

Challenge

OpenX, one of the world's largest independent ad exchanges, needed to replace an unsupported key-value database with a scalable, low-latency alternative. The system had to handle over 150 billion ad requests per day while maintaining sub-10 millisecond response times.

Solution

OpenX migrated to Google Cloud and adopted Bigtable as its primary database for high-volume, low-latency ad serving. Each Bigtable instance is regionally replicated and integrated with Google Kubernetes Engine components that write event data from Pub/Sub. The team later added Memorystore for Memcached to serve 80% of read traffic via cache, further reducing Bigtable load and latency.

Outcomes

- Bigtable handled over **1 million ad requests per second** with P99 latency under 20 milliseconds
- Adding Memorystore **reduced node count by over 50%**
- **Total database infrastructure costs dropped by 50%**

[Read the full story](#)

“With Bigtable and Memorystore, we could leave the problems of our legacy database behind and position ourselves for growth with solutions that provided low latency and high performance in a scalable, managed solution.”

BOGUSŁAW GORCZYCA

Technical Lead, OpenX



Industry
Media and entertainment

Country
Australia

Google Cloud database
[Firestore](#)

Additional Google Cloud products
App Engine

Made in Katana achieves zero-downtime fan voting with Firestore

Challenge

Creative agency Made in Katana (MIK) produces Hottest 100, a wildly popular fan-voted music event in Australia. With each year breaking new audience records, the team needed a database that could scale globally and eliminate DevOps overhead.

Solution

Migrating its voting platform to Firestore's non-relational structure allowed seamless horizontal scaling for MIK, even during real-time spikes from millions of users. App Engine simplified operations with fast spin-up times and eliminated the need for manual infrastructure management. The team restored popular interactive features like shortlist voting and confidently launched new engagement experiences.

Outcomes

- Voting platform achieved **zero downtime** during a record-breaking year with over **3 million fan votes**
- Firestore and App Engine eliminated DevOps burden, **freeing the team to focus on product improvements**
- App Engine **scaled to 100 server instances in under 10 seconds** to meet surging global traffic

[Read the full story](#)

“With Google Cloud, the experience for the end user is less frustrating and a lot more seamless and immediate, resulting in fewer users abandoning the voting process. It’s a key strategic success we’ve gained through a combination of Firestore and App Engine.”

LUKE LARSEN

Senior Web Developer, Made in Katana

Telecom

Adapt quickly, personalize service, and keep your network running strong with intelligent data infrastructure.





Industry
Telecom

Country
United Kingdom

Google Cloud database
[Memorystore for Redis](#)

Additional Google Cloud products
BigQuery, Dataflow

Virgin Media O2 achieves sub-millisecond latency with Memorystore

Challenge

Telecom provider Virgin Media O2 needed to modernize its Netpulse analytics platform, which combines weblog and location data for real-time insights into network performance. Its on-premises Hadoop system couldn't support real-time analytics, suffered from capacity issues, and lacked disaster recovery capabilities.

Solution

Virgin Media O2 rebuilt Netpulse using Memorystore for Redis as a fast, in-memory lookup service to support real-time joins between weblog and mobility data. It selected Redis for its sub-millisecond latency, high availability, and familiar data structures. BigQuery was used for downstream analytics, and Dataflow powered real-time data transformation at scale.

Outcomes

- Infrastructure handles up to **300,000 writes and 1.2 million reads per second**
- Virgin Media O2 **cut infrastructure costs by 66%** by switching from 9 standalone nodes to 20 smaller Redis Cluster shards
- Latency improved to **under 100 microseconds per Redis operation** during peak load

[Read the full story](#)

“With the migration to Memorystore, we’ve laid the foundation for seamless integration of weblogs and MME data in real time. This critical development allows VMO2 to take advantage of a wide range of capabilities that come with this innovative solution.”

CHANDU BHUMAN

Senior Manager, Data Engineering, Virgin Media O2

**Industry**

Telecom

Country

United States

Google Cloud database[Cloud SQL](#)

Broadcom migrates terabyte-scale databases to Cloud SQL

Challenge

Following its acquisition of Symantec's Enterprise Security business, Broadcom needed to migrate 40+ MySQL databases – totaling over 10 TB – to Google Cloud. To meet SLAs for its Symantec Endpoint Security SaaS platform, the team had to minimize downtime to under 10 minutes.

Solution

Broadcom worked with Google Cloud's Professional Services team to implement a continuous data migration strategy using Cloud SQL. The team opted for external server replication instead of a fully managed service to maintain fine-grained control, applying custom stored procedures and schema-level filters. A one-time parallel dump using mydumper/myloader was followed by secure, VPC-peered replication pipelines to maintain data sync. This approach supported complex multi-source, multi-destination scenarios while maintaining security and data fidelity.

Outcomes

- Broadcom migrated 40+ MySQL databases **totaling over 10 TB**
- During migration, **database cutover downtime was under 10 minutes**
- **Automated failover, filtering, and integrity validation** were enabled across all replication flows

[Read the full story](#)

“To handle our complex requirements, we built a highly automated and secure migration pipeline that gave us full control while minimizing disruption.”

RUDRESHA MURTHY

Technical Director, Symantec Endpoint Security Division, Broadcom

Get started with Google Cloud databases

Google Cloud offers the only suite of industry-leading databases built on planet-scale infrastructure with AI at its core.



AlloyDB for PostgreSQL

Run high-performance, AI-powered applications and agents with AlloyDB for PostgreSQL

AlloyDB for PostgreSQL combines the familiarity of PostgreSQL with the best of Google. It delivers high-speed transactional and analytical processing, while eliminating the operational overhead of traditional databases.

With the latest advancements in AlloyDB AI, including enhanced vector processing and natural language support, organizations can power intelligent applications, extract real-time insights, and scale seamlessly to meet demand. Its cost-efficient architecture optimizes compute and storage for high performance without excessive costs. Whether modernizing legacy systems or building AI-driven applications and agents, AlloyDB for PostgreSQL provides the flexibility and efficiency needed for next-generation workloads.

Fast performance

4x faster for transactional workloads and up to 2x better price-performance compared to self-managed PostgreSQL

Built-in AI support

Generate vector embeddings, connect to tools like LangChain, and access models in Vertex AI – all from within your database

Real-time business insights

Leverage its built-in columnar engine, run BI, reporting, and HTAP workloads with up to 100x faster analytical queries, with no impact on operational performance

See how companies are using AlloyDB for PostgreSQL

Retail and consumer packaged goods

Tchibo

Financial services

Galxe
Apex Fintech
FLUIDEFI

Technology

Endear
B4A

Startups and digital natives

Character.AI
Writer.com
Loyal Guru

Healthcare and life sciences

Bayer
NeuroPace


Security

Tricent

Keep exploring AlloyDB for PostgreSQL



AlloyDB for PostgreSQL: the database built for you
Discover how AlloyDB for PostgreSQL combines our love of PostgreSQL with the best of Google.
[Download the ebook](#) →



AlloyDB for PostgreSQL documentation
Find quickstarts and guides, review key references, and get help with common issues.
[Explore documentation](#) →



AlloyDB for PostgreSQL free trial
Get started with a 30-day AlloyDB for PostgreSQL free trial instance.
[Try AlloyDB for PostgreSQL for free](#) →



Simplify database operations with Cloud SQL

Cloud SQL is a fully managed, cost-effective relational service for PostgreSQL, MySQL, and SQL Server built to simplify operations without sacrificing performance or flexibility. It's ideal for lift-and-shift migrations, third-party apps, and new cloud-native builds that need to get up and running fast.

With built-in automation for backups, failover, scaling, and patching, Cloud SQL takes care of database maintenance so teams can focus on innovation. It offers price-performance options for every workload and integrates easily with tools developers already know and love.

Built-in efficiency

Automatic backups, replication, and maintenance reduce manual ops and risk

Optimized price performance

Choose the right capabilities for your performance needs, with options ranging from Enterprise edition and Enterprise Plus edition to new C4 instances powered by Google Axion processors

High availability with near sub-second downtime maintenance

Effortlessly enable high availability with automatic failover and achieve near-zero downtime during maintenance and instance scale-up

See how companies are using Cloud SQL

Financial services

Deckmatch
Sanitas

Technology

Google Nest
Auto Trader
Visual Research
Cart.com

Startups and digital natives

Linear
Lightricks

Healthcare and life sciences

Intelligencia AI

Games

Chess.com

Manufacturing and supply chain

Ford
Renault Group
Manhattan Associates


Security

NetRise

Telecom

Broadcom


Keep exploring Cloud SQL



Resiliency with Cloud SQL

Learn how migrating your databases to Cloud SQL can lower costs, boost agility, and speed up deployments.


[Read the report](#) →



Cloud SQL documentation

Find quickstarts and guides, review key references, and get help with common issues.

[Explore documentation](#) →



Start your proof of concept

New customers get \$300 in free credits to try Cloud SQL and other Google Cloud products.

[Get started for free](#) →



Power global apps that never go offline with Spanner

Spanner is built for availability, scale, and simplicity. With 99.999% availability and horizontal scalability, it keeps mission-critical applications running – even during traffic spikes or global expansion. Spanner’s multi-model capabilities empower you to build intelligent, AI-enabled applications on top of your operational relational and NoSQL data by using native Vertex AI integration, Spanner Graph for querying complex relationships, vector search for semantic search, and built-in full-text search – all with “true ZeroETL” interoperability.

And with Data Boost, you can experience consistently high performance with workload-isolated query processing, even during peak demand.

Global scale

Handles massive read and write traffic without downtime or redesign

Built-in intelligence

Combines relational queries with graph and vector search for richer insights

Always-on availability

Spanner delivers up to 99.999% availability with automated maintenance and flexible deployment options

See how companies are using Spanner

Retail and consumer packaged goods

- Macy’s
- Wayfair
- Kroger
- REWE Group

Financial services

- Deutsche Bank
- CERC
- Current

Technology

- Yahoo

Startups and digital natives

- Uber
- Vimeo
- Prefab
- Kochava

Games

- COLOPL

Manufacturing and supply chain

- Mahindra

Media and entertainment


- Glance

Keep exploring Spanner




The Total Economic Impact™ of Spanner
Discover cost savings and business benefits enabled by Spanner.

[Read the study](#) →



Spanner documentation
Find quickstarts and guides, review key references, and get help with common issues.

[Explore documentation](#) →



Spanner free trial
Create a 90-day Spanner instance for free.

[Try Spanner for free](#) →



Scale low-latency workloads with confidence using Bigtable

Bigtable is Google Cloud's high-throughput, low-latency NoSQL database built for applications that demand speed at scale. It handles everything, from real-time personalization and event tracking to time series and operational analytics – without breaking a sweat.

With native support for wide-column and key-value data, Bigtable can serve, stream, and analyze massive volumes of structured and unstructured data in one place. It scales up or out across regions with automatic sharding and replication, keeping performance consistent and costs predictable.

Real-time performance

Optimized for low-latency reads and writes, even at global scale

Flexible data model

Store JSON, key-value pairs, embeddings, and more in a single database

Built to scale

Add nodes as needed for high throughput – no redesign required

See how companies are using Bigtable

Retail and consumer packaged goods

Macy's
Flipkart

Financial services

Equifax
Symphony

Technology

Sabre
PLAID

Startups and digital natives

Box
Bitly
Airship
Moloco

Manufacturing and supply chain

Ford Pro
Oden Technologies
Bharat Light & Power
STL

Security

Palo Alto Networks
Stairwell
Ravelin

Media and entertainment


YouTube
Grupo Globo
OpenX

Keep exploring Bigtable




SQL support for Bigtable
Bigtable transforms the developer experience with SQL support.

[Learn more →](#)



Bigtable documentation
Find quickstarts and guides, review key references, and get help with common issues.

[Explore documentation →](#)



Take the next step
Start building on Google Cloud with \$300 in free credits.

[Try Bigtable for free →](#)

Memorystore

Move your cache to the cloud with zero complexity using Memorystore

Memorystore brings the speed of in-memory caching to the cloud – no re-architecture, no rewrites. It supports Valkey, Redis* OSS Cluster, Redis OSS, and Memcached, and is fully compatible with open source protocols, so teams can migrate existing deployments as-is.

Whether you're powering a real-time leaderboard, accelerating read-heavy workloads, or supporting AI use cases with vector search, Memorystore delivers sub-millisecond performance and scales easily to meet demand. With built-in availability and failover, your cache stays online when it matters most.

Highly available

Memorystore for Valkey and Redis Cluster supports zero-downtime scaling, automated failover, and a 99.99% SLA (for Redis Cluster)

Built-in vector search

Accelerate generative AI applications with ultra-low latency approximate and exact nearest neighbor vector search

Zero-downtime scaling

Scale Memorystore for Valkey and Memorystore for Redis Cluster up to 250 nodes and 10+ TB per instance with automated failover and 99.99% availability – no interruptions during maintenance or growth

*Redis is a trademark of Redis Ltd. All rights therein are reserved to Redis Ltd. Any use by Google is for referential purposes only and does not indicate any sponsorship, endorsement, or affiliation between Redis and Google. Memorystore is based on and is compatible with open-source Redis versions 7.2 and earlier and supports a subset of the total Redis command library.

See how companies are using Memorystore

Retail and consumer packaged goods
Instacart

Technology
Unity Ads

Startups and digital natives
Character.ai
Statsig

Media and entertainment
FanCode

Telecom
Virgin Media

Keep exploring Memorystore




Introducing Memorystore for Valkey
Memorystore for Valkey is now generally available.

[Learn more](#) →



Memorystore documentation
Find quickstarts and guides, review key references, and get help with common issues.

[Explore documentation](#) →



Take the next step
Start building on Google Cloud with \$300 in free credits.

[Try Memorystore for free](#) →

Firestore

Build fast, dynamic apps with Firestore's serverless document database

Firestore makes it easy to develop rich, responsive applications without managing infrastructure. This fully managed NoSQL database is designed for scale, offering live data sync, offline support, and flexible queries in a JSON-compatible format – and now, MongoDB compatibility, so you can use existing tools and integrations to build faster.

Whether you're building real-time dashboards, collaborative tools, or gen AI apps with vector search, Firestore handles availability, replication, and scaling behind the scenes so you can focus on delivering features – not maintaining databases.

Built for developers

Live sync, offline mode, and rich query support across mobile, web, and server

AI ready

Easily build generative AI applications with Firestore vector search, LangChain, and Llamaindex

Serverless by design

No maintenance, partitioning, or downtime – just seamless, global scale

See how companies are using Firestore

Technology

HighLevel
Endear
B4A

Startups and digital natives

Loyal Guru

Manufacturing and supply chain

Ford

Media and entertainment


Forbes
Made in Katana (MIK)

Keep exploring Firestore




Announcing MongoDB compatibility
Get a complete overview of Firestore's new MongoDB capabilities.

[Read the blog](#) →



Firestore documentation
Find quickstarts and guides, review key references, and get help with common issues.

[Explore documentation](#) →



Deploy and run a dynamic web app
Try an interactive solution that uses Firestore to run a sample application built with JavaScript.

[Sign up to deploy a dynamic website](#) →

**What
will you
build next?**





Take the next step

Ready to see how Brilyant and Google Cloud can transform your business? Our team can help you explore solutions tailored to your specific needs and guide you through implementation. Let's discuss how you can leverage Google Cloud's powerful tools to improve efficiency, drive innovation, and achieve your business goals.

[Contact us](#) today for a personalized consultation and discover how we can help you unlock new possibilities with Google Cloud.

We're here to help you every step of the way.

Sincerely,

Brilyant IT Solutions



Google Cloud
Partner