


CASE STUDY

Why Cisco AI runs on CockroachDB instead of PostgreSQL





Why Cisco AI runs on
CockroachDB instead of
PostgreSQL

INDUSTRY:

AI, Cybersecurity

CHALLENGES:

Meeting AI product requirements for widespread global scale and availability.

SOLUTION:

A modern, scalable, and resilient infrastructure built for AI workloads based on CockroachDB.

About Cisco

[Cisco](#) is a global leader in networking solutions offering a wide range of products to help their customers secure their enterprise, transform their infrastructures, and thrive in an AI-driven world. A few years ago, the company acquired Armorblox, a cybersecurity company that used natural language understanding and AI to secure communications with a focus on email.

Following the acquisition in 2023, the team at Armorblox joined Cisco's Security Business Group and were tasked with adding genAI-powered experiences to the company's security portfolio. Not only did Cisco want to bring AI to their customers, but they also wanted to ensure a safe and secure adoption of AI within their customer's ecosystems.

Armorblox's legacy infrastructure was already creating challenges and they needed to prepare for unprecedented scale. Prior to the acquisition, they decided to migrate from PostgreSQL to CockroachDB, and over the past few years Armorblox (now a part of Cisco) used the globally distributed SQL database to build all new [Cisco AI products](#).

Watch : [Cisco AI runs on CockroachDB](#) ▶

7

REGIONS

65k

TRANSACTIONS PER
SECOND

2M

MAILBOXES



Operational complexity with PostgreSQL

In 2018, Armorblox started building their platform on a microservices-based architecture and decided to run PostgreSQL in Kubernetes. As their business started to grow over the years, they were protecting close to 2 million mailboxes and the databases simply could not keep up.

They ran into several complications with running PostgreSQL in Kubernetes (pod glitches, resource limitations, noisy neighbors) and were doing a lot of complex DevOps work to help scale the system and maintain availability. They were really concerned about meeting customer SLAs so their teams would wake up around 5am PT just to keep the databases up and running for East Coast clients.

They decided to move PostgreSQL to VMs and hired Percona to help manage and maintain the databases. However, they still faced significant scaling challenges, lacked native active-active deployment, and backup and restore wasn't a smooth process.

To add to the complexity, Armorblox was dealing with sensitive customer data and needed to be careful about how it was stored. Instead of creating multi-tenant tables, they ended up creating multi-tenant schemas to keep data isolated and secure. This meant that each tenant had their own schema, and the number of tenants started to grow towards tens of thousands. This volume created new challenges around generating enough transaction IDs and vacuuming tables (500K+).

The engineering team was spending entirely too much money and time on database maintenance. When Cisco approached Armorblox about an acquisition, they knew scale was going to become even more important and they wanted to solve for operational complexity as soon as possible.

Their requirements for a new database solution included:

- ✓ A geographically distributed SQL store
- ✓ Simplified scale
- ✓ Built-in backup and recovery
- ✓ Cloud-managed solution with option to self-host
- ✓ Minimal operational overhead

Another important factor was the team behind the new solution. They wanted to be working with people they can trust to help them solve challenges and provide ongoing support. They started searching for a new solution and came across CockroachDB. The team started reading through the documentation, experimenting with the product (scale/performance tests), and ultimately engaged the Cockroach Labs team.



Scalable and secure AI solutions

After a successful proof-of-concept, Armorblox migrated their email security application to CockroachDB. This was a very write-heavy workload that would see up to 65K transactions per second (which translates to about a half a million operations per second) during peak hours.

The team reported that the transition to CockroachDB was pretty “simple and straightforward.” They basically connected their entire instance to a CockroachDB Cloud instance and started the transition under the guidance of the Cockroach Labs team.

One significant change they made as part of the switch was to migrate to multi-tenant tables as opposed to having a schema per tenant. This way Cisco AI’s customers could still feel like they have their own [dedicated cluster that is isolated](#) from other tenants in terms of performance and security.

Following Cisco’s acquisition of Armorblox, the newly formed Cisco AI team built all of their AI products from scratch on top of CockroachDB including:

- ✓ [Cisco AI Defense](#): Platform that secures AI apps, agents, MCPs and LLMs
- ✓ [Cisco AI Canvas](#): AI orchestrated generative UI interface that connects into all Cisco products (with a focus on network operations and troubleshooting using AI)
- ✓ [Cisco AI Assistant](#): Platform for building AI assistance for all Cisco products

These products leverage a rich dataset to power AI models, chatroom interactions, and automation workflows. Although the products are all still relatively new, the Cisco AI team feels confident knowing that CockroachDB will be ready to scale to support their future needs.

“CockroachDB and the Cockroach Labs team is a seamless fit to help us scale our business and our products. Data is essential for AI, right? Without data, you can’t build better models. Having a solid data platform and operational data layer where we can build better models and better products for our customers is ideal. And that’s what CockroachDB is enabling us to do today.”



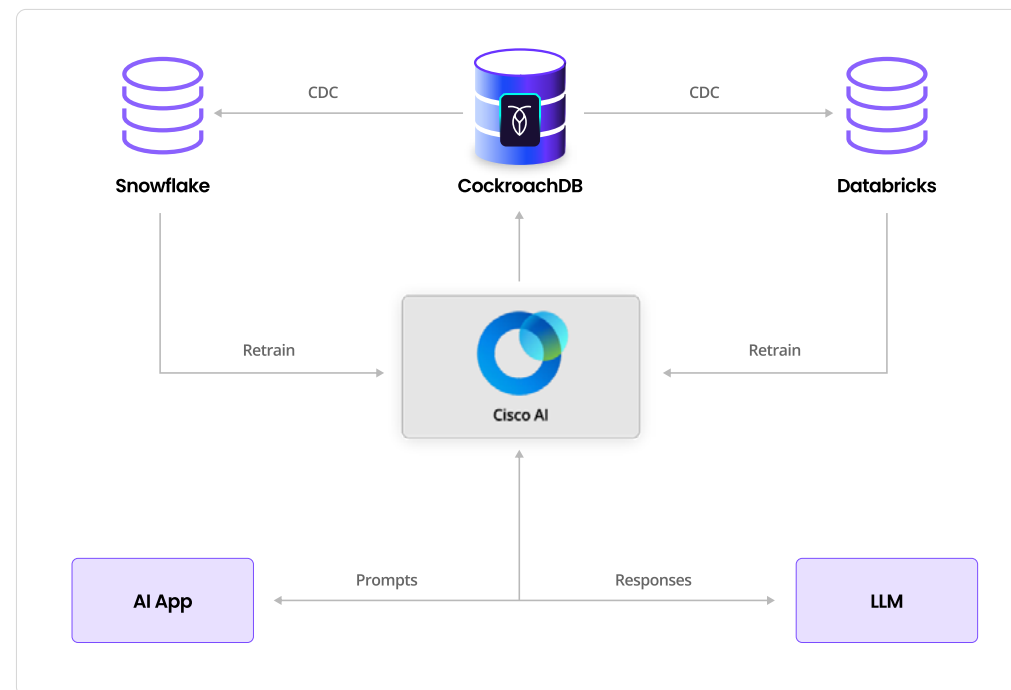
Arjun Sambamoorthy

Sr. Director, Cisco AI

Building a global AI architecture

Today Cisco AI is running CockroachDB across seven different regions across the globe in Europe, Australia, India, and North America. CockroachDB functions as a single logical database regardless of how many regions Cisco AI runs in. This avoids the hassle of maintaining separate clusters per data center, which would double costs and create a need for manual intervention.

Cisco AI's multi-region implementation of CockroachDB enables real-time AI-driven automation which improves efficiency and scalability, and guarantees seamless data synchronization. At a high-level, here's a simplified view of their architecture:



Thousands of Cisco customers interact with the Cisco AI products on a daily basis. All products store prompts from the AI application and responses from the LLM. These responses, along with the entire conversation history, are stored in CockroachDB.

Cisco AI has [change data capture \(CDC\)](#) enabled between CockroachDB and Snowflake so that they can retrain AI safety and security models in the Snowflake pipeline. They also have CDC enabled with Databricks so that they can use the Databricks pipeline to retrain the models used in AI Canvas and AI Assistant products.

“After switching from PostgreSQL to CockroachDB, we don’t spend a lot of time on databases anymore. That’s a huge benefit for our team so we can now focus on what we do best, which is building AI products for Cisco.”



Arjun Sambamoorthy

Sr. Director, Cisco AI

What's next

Cisco AI has been very happy with the results they've seen with CockroachDB. They report that they get great visibility into their systems which helps identify problems and optimize their databases. They've also been pleased with CockroachDB's native scaling capabilities, CDC, and built-in resilience.

As Cisco AI products start to see widespread market adoption, CockroachDB will be the foundation to support the company's growth. In the immediate future, the company wants to expand its presence to South America and grow its customer base in the U.S. and Canada.

To learn more about Cisco AI's product offerings, [click here](#).

Ready to get started?

Go hands-on with 100% free CockroachDB Serverless. Spin up your first cluster in just a few clicks.