

The future of PostgreSQL is here.

And it's built for you.

AlloyDB: The new way to PostgreSQL

Hit the limits of PostgreSQL?

`WARNING: Scaling an app can lead to frustration, burnout, and career-ending rage episodes (!!!)`
Type `"help"` for help.
Or, try `[something new]`.

We get it. It's frustrating to think that your database is holding you back. After all, you can change your own code, but you can't change the database... or can you?

Like many of you, we love building with PostgreSQL — there's a reason it's such a popular open-source database. So, we took our love for PostgreSQL and supercharged it with the best of Google to build [AlloyDB for PostgreSQL](#). AlloyDB is our 100% PostgreSQL compatible database with built-in gen AI capabilities, that runs anywhere. Every aspect is designed for developers like you — to help with your day-to-day role and to drive your ambition forward.

In this e-book, we'll explore the key features that redefine what's possible with PostgreSQL.

- 1 Faster**
- 2 Better**
- 3 Stronger**
- 4 Smarter**
- 5 Cheat sheet**
- 6 What's next?**

AlloyDB is more than 4x faster than standard PostgreSQL for transactional workloads.

TL;DR

Need superior speed without sacrificing accuracy and cost? AlloyDB has you covered.

```
SELECT name, transactions_per_minute  
FROM database_stats  
ORDER BY transactions_per_minute  
LIMIT 2
```

> name	transactions_per_minute
> -----	-----
> AlloyDB 16vCPU	614,797 txn per minute
> Postgres 16vCPU	140,942 txn per minute

Wish your PostgreSQL was 4x faster?

Whisper it: being a developer can be, err, *challenging*.

In 2024, every industry — from e-commerce to finance to gaming — is reporting huge demands on workloads. For developers, this means the stakes are high and the pressure is on. Of course, speed alone doesn't make a database — which is why we built AlloyDB to be not just faster, but more reliable and efficient when working at scale. It's about giving you the tools for smoother, less stressful processing of high-volume transactions so you can actually enjoy building your apps.

At AlloyDB's core is a [storage-optimized database engine](#) that's 100% PostgreSQL compatible. Because we had access to the same Google infrastructure that runs popular services such as YouTube, Search, Maps, and Gmail, we were able to think big and create a database that disaggregates compute and storage at every layer of the stack. AlloyDB also scales horizontally with load-balanced, low-latency read pool instances. Never again will you have to choose between scale, reliability, and cost — instead, you can enjoy a database that's 4x faster than standard PostgreSQL for transactional workloads.

“ We knew that Google Cloud databases offer exceptional scalability and performance, but we didn't expect the superior performance and fast analytical capabilities we get from AlloyDB. We rely on AlloyDB to power all of our core transactional processes, including user profiles, loyalty program details, and offer redemptions. It ensures our APIs have reliable, low-latency access to this critical data. ”



JESÚS ANTONIO CANALES DIEZ
PLATFORM TECH LEAD, LOYAL GURU

Should reporting take 2 hours or 2 minutes?

Launch days, monthly reporting, business insights: database professionals are the natural gatekeepers of crucial organizational information. Which is fine, until you're on a deadline and the distracting requests start pouring in.

With AlloyDB's built-in [columnar engine](#), you can run analytical queries up to 100x faster than standard PostgreSQL. This isn't just about speeding up — we rigorously ensured there will be zero impact on operational performance when running business intelligence, reporting, and hybrid transactional and analytical workloads (HTAP).

You can also stream data or use federated queries with [BigQuery](#) and call machine learning models in [Vertex AI](#) directly within a query or transaction.

AlloyDB provides up to 2x better price-performance compared to self-managed PostgreSQL.

Do you like the sound of AlloyDB but are wondering about the costs? You can show this at the next budget meeting: AlloyDB provides up to 2x better price-performance compared to self-managed PostgreSQL.

Developer insight


“4x faster means even my dodgy code is fast enough!”



A database that slows you down?

Or one that accelerates your success?

AlloyDB Omni is more than 2x faster than standard PostgreSQL for transactional workloads.



TL;DR

With a downloadable version, AlloyDB can tackle any project, anywhere.



```
docker run --name my-omni \  
-e POSTGRES_PASSWORD=NEW_PASSWORD \  
-d google/alloydbomni
```

A cloud database that holds you back? Or one that runs anywhere?

Running your apps in the cloud is great, sure. But some workloads need to run on site - you can't move an oil rig, a telco facility, or a retail store into the cloud. So flexibility has become the cornerstone of any successful database.

With our downloadable edition of AlloyDB, called [AlloyDB Omni](#), you can run your database in the middle of the Mojave desert... if that's what your business needs. This isn't a scaled-back version either. It's powered by the same engine that drives the cloud-based AlloyDB service — so it offers all of the same functionality. And because it runs anywhere, you can easily build apps in virtually any environment whether in the cloud or on-premises. The possibilities are endless!

You work around the world. Shouldn't your database?

We've also launched the AlloyDB Omni Kubernetes operator, which simplifies common database tasks such as database provisioning, backups, high availability, secure connectivity, and observability, allowing you to run AlloyDB Omni in most Kubernetes environments.

AlloyDB Omni also offers a shortcut for modernizing legacy databases, anywhere. At a fraction of the cost, it can be used for getting older databases up to an enterprise-grade version of PostgreSQL supported by a Tier 1 vendor. This means goodbye to the expensive fees, unfriendly licensing relationships, and vendor lock-in associated with legacy database vendors.

Developer insight

"Knowing that I can host my cloud database on-premises, keeping the data within the boundaries of my on-prem environment was game changing!"

“ AlloyDB AI on AlloyDB Omni's indexing and embedding query capabilities streamlines medical analysis, eliminating the need for external processing among millions of patient records. ”




SHARANYA DESAI
PH.D, NEUROPACE



Can you download your cloud database and run it anywhere in the world?

AlloyDB detects and recovers from system failures in 60 seconds or less.



TL;DR

A fully managed database with 99.99% availability SLA, an architecture built for resilience, and auto-pilot features that offer higher reliability as well as ease of use, meaning your work is safe with us.

```
pg_dump  
> pg_dump: Command not required,  
AlloyDB does all backups automatically.
```

Starting to feel more DevOps than Dev?

As part of our vision for a truly developer-friendly database, we knew AlloyDB had to be fully and intelligently managed, as well as offering near-perfect uptime. We also ensured AlloyDB offers refreshingly predictable performance and a high availability service-level agreement (SLA) of 99.99%, inclusive of maintenance.

To make all this happen, AlloyDB uses adaptive, intelligent algorithms to do the work that humans hate to do. PostgreSQL vacuum management? Storage and memory management? Data tiering? Analytics acceleration? Consider it done. Oh, and if you happen to hate query optimization, check out the [Index Advisor](#), which monitors queries and recommends indexes that can be added to increase performance.

We know how every minute matters when you're building the future — or even just trying to get home on time. Which is why AlloyDB detects and recovers from most system failures within 60 seconds and offers a high SLA of 99.99%, inclusive of maintenance. That's your stress-free coffee break secured. And if you're trying to scale, you really can't afford to lose time to slowdowns and outages.

“ We came to appreciate the ease of maintenance that a fully managed solution like AlloyDB provides. In a startup like ours, resources are limited, so the ease of use and integration plus the 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 RIEHLE
CEO AND FOUNDER, B4A

Does your backup have your back?

With point-in-time recovery to any point within your defined retention window, you can restore an AlloyDB database to a specific date and time for development, test and auditing purposes, or recover your production database from user* or application** errors.

*It does happen.

**It happens (not to you, of course).

In addition, AlloyDB can create and store complete backups of your data, either on demand or on a regular schedule. And for disaster recovery, it offers cross-region replication — AlloyDB will stream data from your primary region to one or more secondary Google Cloud regions, and you can promote a database in a secondary region if disaster strikes.

As for the business gain? Well, you can hardly put a price on protecting months, or even years, of work.

Developer insight

“This sounds a bit extra, but not having to worry about my backup scheduler running properly helps me sleep at night.”



Hours spent managing your database?

Or time to actually develop your apps?

AlloyDB AI offers up to 4x faster vector queries, 8x faster index creation, and uses 3-4x less memory than the HNSW index in standard PostgreSQL.



TL;DR

AlloyDB is both the best database for building generative AI applications *and* is equipped with its own impressive AI capabilities.

```
-- create embeddings in AlloyDB from our
existing Postgres chat history table
CREATE TABLE alloydb.embeddings AS
SELECT embedding('textembedding-gecko-
multilingual@001', message), message
FROM postgres.chat_history
```

Wish you were already building the next big gen AI app?

Are your customers suddenly all about gen AI? Or is it finally your turn to get rich quick? (No judgment here). AlloyDB is not only the best database for building gen AI apps, it's also powered by our own smart AI.

Through [AlloyDB AI](#), which is an integral part of the database, you get vector search from a company that knows a thing or two about search. Our new pgvector-compatible index is based on Google's state-of-the-art approximate nearest neighbor algorithms, [the ScaNN index](#). AlloyDB AI offers up to 4x faster vector queries, up to 8x faster index creation, and typically uses 3-4x less memory than the HNSW index in standard PostgreSQL.

AlloyDB AI offers a suite of tools for easily building enterprise gen AI apps, while also addressing some of the memory, indexing speed, and querying performance requirements encountered in demanding real-world use cases. You can also generate embeddings from within your database, and we fully integrate with [Vertex AI](#) and the [LangChain](#) open source framework.

Developer insight

"Being able to use my existing database structure for a new AI project is an absolute game changer for me."

“AlloyDB AI acts as a dynamic vector store, indexing repositories of regulatory guidelines, compliance documents, and historical reporting data to ground the chatbot. Compliance analysts and reporting specialists interact with the chatbot in a conversational manner, saving time and addressing diverse regulatory reporting questions.”



ANTOINE MOREAU
CIO, REGNOLOGY

What could Gemini do for you?

Part of a new generation of AI foundation models, Gemini gives you a way to supercharge your ideas and completely change the way you work.

If you've found yourself bored, chipping away at the same repetitive SQL coding tasks or trying to decipher database performance dashboards, you should give Gemini a try. It supports our vision for a database that lets you build without limits — just imagine what you could do with those repetitive day-to-day frustrations taken care of.

Here are just a few of the Gemini-assisted capabilities tools we offer through AlloyDB:

- **Database Studio** offers AI-assisted capabilities to easily generate SQL code using natural language. Intelligent code assistance and code completion are available directly within the Google Cloud console.
- **Database Center** simplifies all aspects of database operations including performance optimization, fleet management, and governance. It offers insights into database performance, security, reliability, and cost.
- **Database Migration Service** offers code assistance and schema conversion, in preview. It revolutionizes database migrations with LLM-powered conversion for the last mile of migration.



Will you stand on the sidelines?

Or take advantage of AlloyDB's innovation?

↳ Need something to show your organization? Here's the cheat sheet.

Better performance and scale for PostgreSQL workloads.

Better price-performance than standard PostgreSQL.

Super easy management.

Easy development of generative AI and semantic search applications.

Runs anywhere — in any cloud, in your data center, even on your laptop.

Backed by a Tier 1 cloud provider.

Provides fast, real-time business insights with no impact on operational database performance.

Embedded Gemini assistive capabilities to support app development, database administration, and migration.

/So, does your existing database stack up? *



Now you've seen what AlloyDB can do, and more importantly, what **you** could do with AlloyDB, you might have a few questions.

For expert advice from a member of the Cloud team, [contact us](#).

Or, try AlloyDB for free today: goo.gle/try_alloydb