

Tembo Support Docs

Tembo Cloud

Tembo is a dev-first, fully-extensible, fully-managed, secure, and scalable Postgres service. Tembo provides the largest library of easily-installed extensions and use-case optimized **Tembo Stacks**, allowing our customers to use Postgres for all their data needs.

For more, login to cloud.tembo.io

The screenshot shows the 'Active Instances' page in the Tembo Cloud interface. At the top, there is a navigation bar with the Tembo logo, 'tembo', 'My Instances', 'Docs', and 'Support'. A 'Create an instance' button is located in the top right. The main content area displays four instance cards, each with a 'GOOD' status indicator.

Active Instances

Create an instance

| Instance Name | Stack | Status |
|--------------------|--------------------|--------|
| control-plane-prod | Standard Stack | GOOD |
| cp-queue-prd | MessageQueue Stack | GOOD |
| pgtrunkio-prod | Standard Stack | GOOD |
| tembo-dw-prod | OLAP Stack | GOOD |

| Extensions | Provider | Region | CPU | RAM | Storage |
|------------|----------|-----------|-------|-----|---------|
| 10 | AWS | US East 1 | 1vCPU | 4Gi | 10Gi |

| Extensions | Provider | Region | CPU | RAM | Storage |
|------------|----------|-----------|-------|-----|---------|
| 10 | AWS | US East 1 | 1vCPU | 4Gi | 10Gi |

| Extensions | Provider | Region | CPU | RAM | Storage |
|------------|----------|-----------|-------|-----|---------|
| 9 | AWS | US East 1 | 1vCPU | 1Gi | 10Gi |

| Extensions | Provider | Region | CPU | RAM | Storage |
|------------|----------|-----------|-------|-----|---------|
| 19 | AWS | US East 1 | 2vCPU | 4Gi | 100Gi |

Tembo Stacks

“Postgres for Everything” delivered as polished “flavored” Postgres — Tembo Stacks. We help teams avoid introducing new databases, and the associated pains. Less database sprawl.

- Tembo OLTP
 - A low latency transactional processing stack with high I/O performance, concurrency, and real-time metrics.
- Tembo Vector
 - An alternative to Pinecone, Weaviate, and Chroma, built on Postgres. Store and search vectors with Pgvector. Manage the lifecycle of embeddings with pg_vectorize.
- Tembo Timeseries
 - Easily build reliable and performant time-series processing on PostgreSQL.
- Tembo Messaging Queue
 - An alternative to AWS SQS and Redis RSMQ, on Postgres. Interface to Postgres queues using SQL, REST API, and many client libraries including Python, Rust, and Go.
- Tembo Geospatial
 - Postgres pre-configured for Geospatial workloads. PostGIS comes pre-installed.
- Tembo DataWarehouse
 - Extract, Transform and Load data from external sources. Build centralized datastore for analytical and tactical queries.
- Tembo OLAP
 - Postgres tuned for online analytical processing, optimized for large data sets, complex queries, and high throughput.
- Tembo RAG
 - An alternative to LangChain and LlamaIndex. Build LLM applications faster and with less infrastructure on Postgres.
- Mongo Alternative on Postgres
 - Mongo-compatible wire protocol on Postgres. Connect your Mongo client to Postgres with zero code changes.
- Tembo ML
 - Machine learning training and inference directly from Postgres. Built on PostgresML, pg_vector and pg_vectorize.
- Tembo Standard
 - Balanced for general purpose computing with full control over compute, configuration, and extension installation.

Select a Stack

Stacks are specialized Postgres instances that are designed for your workload. Click any stack to learn more.

OLTP
A low latency transactional processing stack with high I/O performance, concurrency, and real-time metrics.
7 extensions: transaction-processing, real-time-operations, operational data

VectorDB
An alternative to Pinecone, Weaviate, and Chroma, built on Postgres. Store and search vectors with pgvector. Manage the lifecycle of embeddings with pg_vectorize.
4 extensions: vector-search, similarity-search, pgvector, embeddings

Time Series
Easily build reliable and performant time-series processing on PostgreSQL.
2 extensions: trend-analysis, data-analysis, real-time, forecasting, dashboarding

Message Queue
An alternative to AWS SQS and Redis RSMQ, on Postgres. Interface to Postgres

OLTP

Overview
The OLTP stack is a finely-tuned database cluster optimized for transactional workloads, designed to handle concurrency with ease. Built with optimized WAL and auto-vacuum settings, it also includes extensions for debugging and real-time metrics.
transaction-processing, real-time-operations, operational data, primary data store, high throughput

Configuration
The following configurations automatically scale based on the size of cpu, memory, and storage for the cluster:
• shared_buffers
• max_connections

Deploy Now

What is a Tembo Stack?

Stacks are pre-built, use case-specific configurations of Postgres, enabling you to quickly deploy specialized data services that can replace external, non-Postgres data services. They include:

- Docker Base Image (Postgres Container)
- Pre-Installed Extensions
- Stack-Specific additional metrics, alerts + recommendations
- Stack-Optimized Postgres Configs
- Dynamic Config Engine
- Sidecars (i.e. nearby workloads)
- Infrastructure and hardware configuration options - things like the HA setup, pgbouncer (server-side connection pooling), etc.