


Building Trunk:

A Postgres Extension
Registry and CLI Supporting
Tembo Cloud

Ian Stanton
Tembo

Backstory



Backstory

Tembo

Founded November, 2022

Managed Postgres solution - Tembo Cloud

Vision

Postgres for everything. Use-case optimized flavors of Postgres, powered by extensions

Ability to install and use any extension, even user-provided extensions

We knew we'd need some extension management system to pull this off

The screenshot displays the Tembo Cloud website, which promotes Postgres extensions. On the left, a vertical list of extension categories is shown, each with a Postgres icon, a description, and a list of supported extensions:

- DataWarehouse**: Extract, Transform and Load data from external sources. Build centralized datastores for analytical and tactical queries. Extensions: 8 extensions, connectors, data-analysis, reporting, etl, fdw.
- OLAP**: Postgres tuned for online analytical processing, optimized for large data sets, complex queries, and high throughput. Extensions: 7 extensions, analytical-processing, data-analysis, business-intelligence.
- Retrieval Augmented Generation**: An alternative to LangChain and Llamaindex. Build LLM applications faster and with less infrastructure on Postgres. Extensions: 3 extensions, rag, generative ai, chat bot, nlp.
- Geospatial**: Postgres for all your geospatial needs. Extensions: 7 extensions, spatial-data, gis, postgis, location-based-services.
- Machine Learning**: Postgres for machine learning training and inference workloads. Extensions: 4 extensions, ml-ops, supervised-learning, postgresml, pgml. (Marked as PREMIUM)
- Message Queue**: Configured for message queue workloads using the PGMQ extension, similar to AWS SQS and Redis Simple Message Queues. Extensions: 5 extensions, asynchronous-communication, message-broker, pgmq, pub-sub-model.
- Mongo Alternative on Postgres**: Mongo-compatible wire protocol on Postgres. Extensions: 1 extensions, nosql, mongodb, ferretdb, document-store, json-storage.

On the right, a blue banner highlights the **Message Queue** feature. Below it, an **Overview** section explains that Message queues allow sending, reading, and retaining messages between applications without data loss, powered by PGMQ (a Postgres extension). It notes that the Message Queue Stack is pre-installed on Postgres clusters and that infrastructure configurations are optimized for message queue workloads. A list of tags includes: asynchronous-communication, message-broker, pgmq, queue-management, event-driven-architecture, and pub-sub-model.

The **PGMQ Features** section lists:

- Lightweight, built with Rust and Postgres only
- Guaranteed "exactly once" delivery of messages consumer within a visibility timeout
- API parity with AWS SQS and RSMQ
- Messages stay in the queue until explicitly deleted
- Messages can be archived, instead of deleted, for long-term retention and replayability
- Table (bloat) maintenance automated with pg_partman
- High performance operations with index-only scans.

A **Credits** section mentions: The PGMQ extension.

Backstory

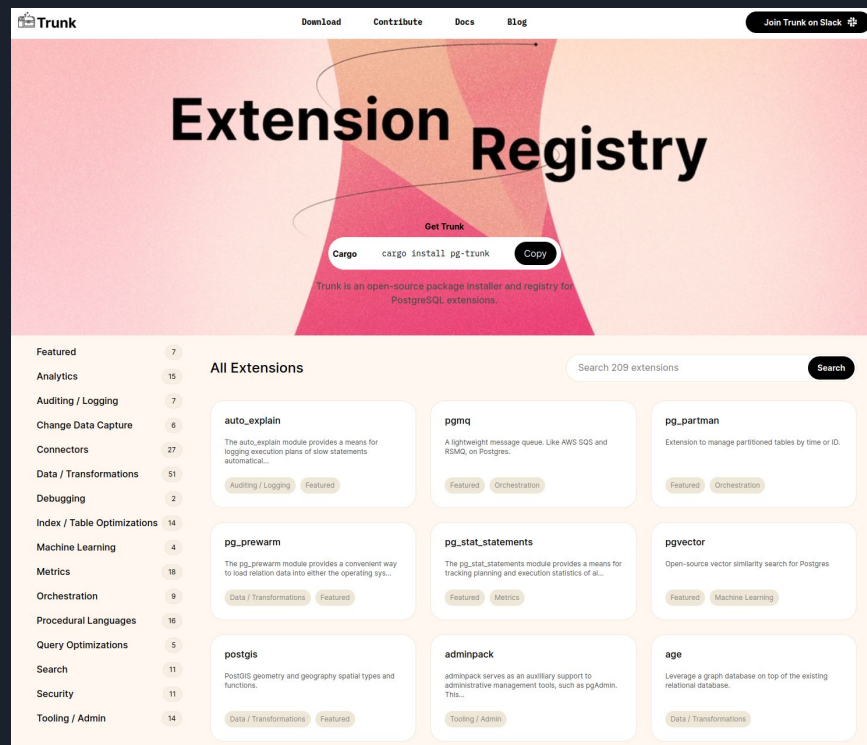
Trunk

A Postgres extension registry and CLI

Binary extension packaging

Open to the community and open source

Powers Tembo Cloud's extension management experience





Backstory

Tembo Cloud

We spent Q1 2023 building Tembo Cloud V1

Extensions were baked into our Postgres container images

We wanted a way install extension binaries on the fly without rebuilding container images

Extension Registry

Started exploring the extension ecosystem for existing solutions



Backstory

PGXN

- + Backed by the community
- + Around since 2011
- Source distribution, not binary
- Little development since 2012


apt / yum

- + Backed by the community
- + Binary distribution
- Smaller subset of extensions available
- Review process for adding could slow us down

Building our own

- + Full control over integration with Tembo Cloud
- + Binary distribution
- + We could build new features quickly
- + We could publish new extensions quickly
- How will the community react?
- Recreating the wheel?

Inspiration



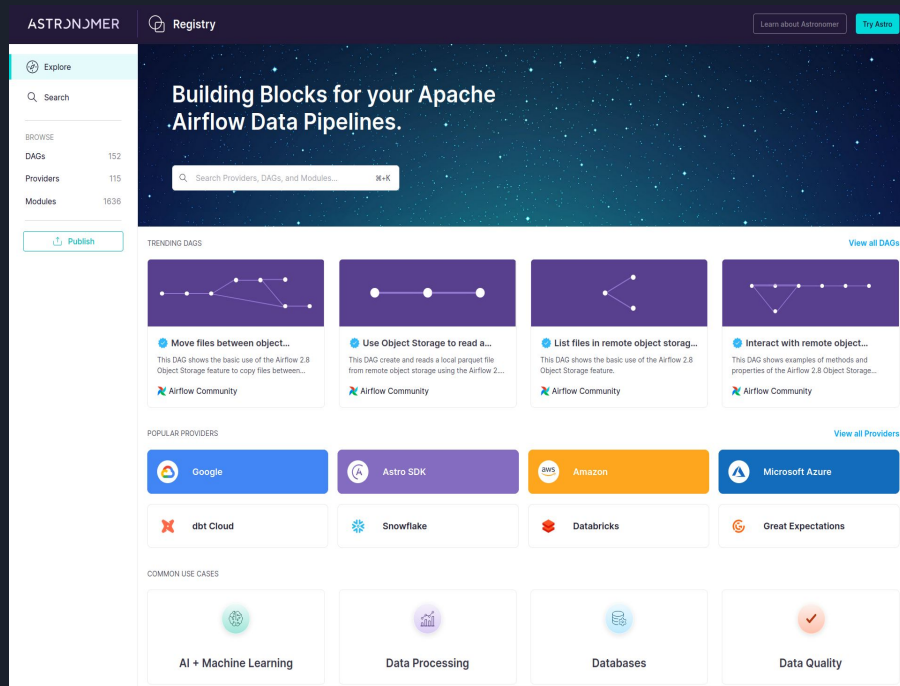
Inspiration

Astronomer Registry for Airflow

A handful of Tembo's employees worked at Astronomer - managed service for Apache Airflow

Community registry for DAGs, Providers and Modules

We wanted to build something similar, but for Postgres extensions





Inspiration

PGXN / apt / yum

We wanted to take the best of what the current offerings had to offer and build upon that


Crates.io

We wanted the experience to be as simple and enjoyable as building, publishing and installing Rust crates.

Vision

- Discoverability
 - Categories
 - Rating system
 - Certification
 - Cloud provider support
- Build any extension
 - SQL
 - C/SQL
 - PGRX
- Simple publishing
 - Like `cargo publish`
- Simple installation
 - Like `cargo install cargo-pgrx`

Hopes and Dreams





Hopes and Dreams

Community involvement


- Extension publications
- Code contributions

Multi-platform support

Multiple versions of Postgres

The one-stop-shop for Postgres extensions

How It Works





How It Works

CLI

Built in Rust

Uses different build images based on extension type detected

Packages extension files into
`<trunk-project-name>-<version>-<pg-version>.tar.gz`

Publishes to registry with `trunk publish`


Installs from registry `trunk install`

Registry

Built in Rust

Backend stores extension metadata

S3 bucket stores `tar.gz` file



Create Trunk bundle:

```
bitcode/src/semver/src/semver.bc
bitcode/src/semver.index.bc
semver.so
licenses/LICENSE
extension/semver--0.10.0--0.11.0.sql
extension/semver--0.11.0--0.12.0.sql
extension/semver--0.12.0--0.13.0.sql
extension/semver--0.13.0--0.15.0.sql
extension/semver--0.15.0--0.16.0.sql
extension/semver--0.16.0--0.17.0.sql
extension/semver--0.17.0--0.20.0.sql
extension/semver--0.2.1--0.2.4.sql
extension/semver--0.2.4--0.3.0.sql
extension/semver--0.20.0--0.21.0.sql
extension/semver--0.21.0--0.22.0.sql
extension/semver--0.22.0--0.30.0.sql
extension/semver--0.3.0--0.4.0.sql
extension/semver--0.30.0--0.31.0.sql
extension/semver--0.31.0--0.31.1.sql
extension/semver--0.31.1--0.31.2.sql
extension/semver--0.31.2--0.32.0.sql
extension/semver--0.32.1.sql
extension/semver--0.5.0--0.10.0.sql
extension/semver--unpackaged--0.2.1.sql
extension/semver.control
extension/semver.sql
manifest.json
```

Packed to ./trunk/pg_semver-0.32.1-pg15.tar.gz

```
▼ 1:
  name: "pg_semver"
  description: "A semantic version data type for PostgreSQL."
  documentation_link: "https://github.com/theory/pg-semver"
  repository_link: "https://github.com/theory/pg-semver"
  version: "0.32.1"
  ▼ postgres_versions:
    0: 15
    1: 14
    2: 16
  ▼ extensions:
    ▼ 0:
      extension_name: "semver"
      version: "0.32.1"
      trunk_project_name: "pg_semver"
      dependencies_extension_names: null
      loadable_libraries: null
      configurations: null
      ▼ control_file:
        absent: false
        ▼ content: "# semver extension\ncomment = 'Semantic version data type'\ndefault_version = '0.32.1'\nmodule_pathname = '$libdir/semver'\nrelocatable = true\n"
  ▼ downloads:
    ▼ 0:
      ▼ link: "https://cdb-plat-usel-prod-pgtrunkio.s3.amazonaws.com/extensions/pg_semver/pg_semver-pg14-0.32.1.tar.gz"
      pg_version: 14
      platform: "linux/amd64"
      ▼ sha256: "f412c7b4722eac32a38dbcc7cd4201d95f07fd88b7abc623cd84c77aecc8d4bb"
    ▼ 1:
      ▼ link: "https://cdb-plat-usel-prod-pgtrunkio.s3.amazonaws.com/extensions/pg_semver/pg_semver-pg15-0.32.1.tar.gz"
      pg_version: 15
      platform: "linux/amd64"
      ▼ sha256: "9213771ffc44fb5a88726770f88fd13e62118b0f861e23271c3eeee427a23be9"
    ▼ 2:
      ▼ link: "https://cdb-plat-usel-prod-pgtrunkio.s3.amazonaws.com/extensions/pg_semver/pg_semver-pg16-0.32.1.tar.gz"
      pg_version: 16
      platform: "linux/amd64"
      ▼ sha256: "8ffe4fa491f13a1764580d274e9f9909af4461aacbeb15857ab2fa235b152117"
```

Search[Download](#)[Contribute](#)[Docs](#)[Blog](#)[Join Trunk on Slack](#)

pg_semver

A semantic version data type for PostgreSQL.

Install

```
trunk install pg_semver
```



semver 0.32.0

pgxn 0.32.1
CI passing

This library contains a single PostgreSQL extension, a data type called "semver". It's an implementation of the version number format specified by the [Semantic Versioning 2.0.0 Specification](#).

Installation

To build semver:

```
make
make install
make installcheck
```

If you encounter an error such as:

Details

About

Category Data / Transformations

Version 0.32.1


Last updated about 2 months ago

License PostgreSQL

Architecture x86-64 ⓘ

Operating system Debian/Ubuntu ⓘ

github.com/theory/pg-semver



```
> trunk install pg_semver
Using pkglibdir: "/usr/lib/postgresql/16/lib"
Using sharedir: "/usr/share/postgresql/16"
Using Postgres version: 16
info: Downloading from: https://cdb-plat-use1-prod-pgtrunkio.s3.amazonaws.com/extensions/pg_semver/pg_semver-pg16-0.32.1.tar.gz
info: Dependent extensions to be installed: []
info: Installing pg_semver 0.32.1
[+] bitcode/src/semver/src/semver.bc => /usr/lib/postgresql/16/lib
[+] bitcode/src/semver.index.bc => /usr/lib/postgresql/16/lib
[+] semver.so => /usr/lib/postgresql/16/lib
info: Skipping license file licenses/LICENSE
[+] extension/semver--0.10.0--0.11.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.11.0--0.12.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.12.0--0.13.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.13.0--0.15.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.15.0--0.16.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.16.0--0.17.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.17.0--0.20.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.2.1--0.2.4.sql => /usr/share/postgresql/16
[+] extension/semver--0.2.4--0.3.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.20.0--0.21.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.21.0--0.22.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.22.0--0.30.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.3.0--0.4.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.30.0--0.31.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.31.0--0.31.1.sql => /usr/share/postgresql/16
[+] extension/semver--0.31.1--0.31.2.sql => /usr/share/postgresql/16
[+] extension/semver--0.31.2--0.32.0.sql => /usr/share/postgresql/16
[+] extension/semver--0.32.1.sql => /usr/share/postgresql/16
[+] extension/semver--0.5.0--0.10.0.sql => /usr/share/postgresql/16
[+] extension/semver--unpackaged--0.2.1.sql => /usr/share/postgresql/16
[+] extension/semver.control => /usr/share/postgresql/16
[+] extension/semver.sql => /usr/share/postgresql/16
```

```
*****
* POST INSTALLATION STEPS *
*****
```


Install the following system-level dependencies:

On systems using apt:
libc6

Enable the extension with:

CREATE EXTENSION IF NOT EXISTS semver CASCADE;

State of the Project



State of the Project

Tembo Cloud

200+ extensions

Powers extension management experience

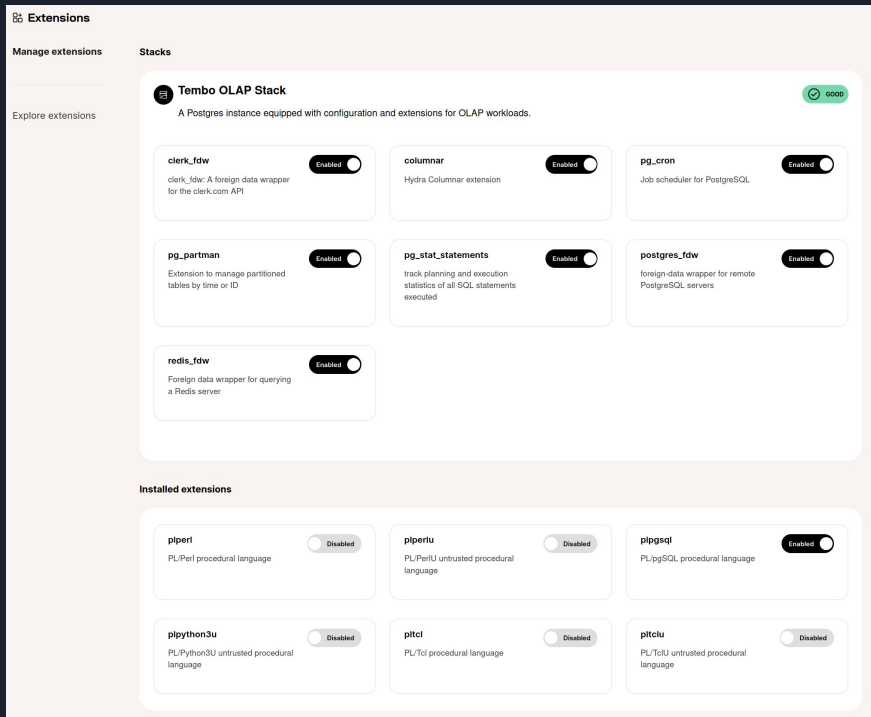
Likely Trunk's #1 user

Community

A lot of site traffic

Minimal community involvement

Repo is not contributor friendly





State of the Project

Did We Solve The Problem?

For Tembo Cloud... yes!

For the community... no!

But we helped bring more awareness to opportunities for improving the extension ecosystem

And more extension management tools were introduced

Why we're all here today!

The Database Package Manager

For Postgres Trusted Language Extensions

Getting started

Documentation

Download and install

Install extensions via a SQL command

```
select dbdev.install('olrice', 'index_advisor');
```

PUBLISHER HANDLE

EXTENSION NAME

> extensions, made easy


npm for PostgreSQL

Brew

Curl

```
brew install pgxman/tap/pgxman
```

Lessons Learned



Lessons Learned

It's a really hard problem to solve!

The extension ecosystem is the wild west!

Community involvement is crucial!

We should have reached out to David in the very beginning!



The Future of Trunk





The Future of Trunk

Trunk will continue to be a core component of the extension management solution in Tembo Cloud

It's paved the way for binary packaging and distribution in PGXN v2

Will inform and be the basis of binary packaging / distribution / installation features in PGXN v2

Thank You!

