

EDB® TECH GUIDE

6 THINGS A DEVELOPER SHOULD KNOW ABOUT POSTGRES



EnterpriseDB

The Enterprise Postgres Company
www.enterprisedb.com

EDBTM
POSTGRES

CONTENTS

03

Introduction

05

No Need to Learn a New IDE

06

Simple Local Setup and Configuration

07

Flexible Multi-Model Architecture

08

Agile and DevOps Friendly

10


Cool Things You Can Do With PostgreSQL

12

Smooth Road to Production

13

Final Note





6 THINGS A DEVELOPER SHOULD KNOW ABOUT POSTGRES

INTRODUCTION

For a long time now, PostgreSQL has been one of the most frequently deployed relational database management systems. Over the past few years, however, it has seen a sharp rise in popularity with developers. According to [DB-Engines](#), PostgreSQL (or Postgres) was ranked the fourth most popular database management system in April 2019, as scored by search engine results, technical discussions, social and professional network mentions, and job offers:

345 systems in ranking, April 2019


Rank			DBMS	Database Model	Score		
Apr 2019	Mar 2019	Apr 2018			Apr 2019	Mar 2019	Apr 2018
1.	1.	1.	Oracle +	Relational, Multi-model	1279.94	+0.80	-9.85
2.	2.	2.	MySQL +	Relational, Multi-model	1215.14	+16.89	-11.26
3.	3.	3.	Microsoft SQL Server +	Relational, Multi-model	1059.96	+12.11	-35.55
4.	4.	4.	PostgreSQL +	Relational, Multi-model	478.72	+8.91	+83.25
5.	5.	5.	MongoDB +	Document	401.98	+0.64	+60.57
6.	6.	6.	IBM Db2 +	Relational, Multi-model	176.05	-1.15	-12.89
7.	↑8.	↑9.	Redis +	Key-value, Multi-model	146.38	+0.25	+16.27
8.	↑9.	8.	Elasticsearch +	Search engine, Multi-model	146.00	+3.21	+14.64
9.	↓7.	↓7.	Microsoft Access	Relational	144.65	-1.55	+12.43
10.	10.	↑11.	SQLite +	Relational	124.21	-0.66	+8.23

Source: <https://db-engines.com/en/ranking>, April 2019

We also learn from the [2018 Stack Overflow survey](#) of more than 100,000 developers that PostgreSQL ranked as the third most used database technology:

- MySQL: 58.7%
- SQL Server: 41.2%
- PostgreSQL: 32.9%
- MongoDB: 25.9%
- SQLite: 19.7%

In short, Postgres has come a long way since its origins as a research project in the mid-1980s at the University of California, Berkeley. Its first step towards the mainstream was in 1995 when it replaced its query language (PostQUEL) with SQL, the de facto industry standard. The very next year, PostgreSQL left academia with the establishment of the PostgreSQL Global Development Group, which took the open-source Postgres codebase under its wings. Since the release in December 2014 of PostgreSQL 9.4 (the oldest of the versions currently supported), its popularity has risen steadily, as shown by [DB-Engines historical graph](#):



There are a number of factors that have contributed to PostgreSQL's popularity, not the least of which is its highly active open-source community that, unlike a company-led open source DBMS like MongoDB or MySQL, is not controlled by any single sponsor or company. Its feature set is continually being extended, with the community testing each new feature thoroughly before integration, so that the highest level of reliability is upheld. Despite its wealth of features, Postgres is lightweight (the source code is less than 20 MB), easy to install, and easy to move around compared to bulky proprietary RDBMS. It is also highly flexible, supporting relational and non-relational models in the same DB, and can be deployed across a wide range of infrastructures from bare metal to VMs, containers, and DBaaS.

In this tech guide, we describe six things that a developer should know about in terms of how easy it is to install, work with, and deploy Postgres as well as the unique capabilities and wide range of features that Postgres offers.

1. No Need to Learn a New IDE

When you opt for a proprietary DBMS, you pretty much have to buy into their development and administrative tooling as well. If you choose to work with PostgreSQL, however, it is unlikely that you will need to install and learn a new IDE, because chances are that the framework you are used to working with is already integrated with PostgreSQL.


[The Community Guide to PostgreSQL GUI Tools on Wiki](#) currently lists no less than 18 freeware and 63 proprietary IDEs that support PostgreSQL, including:

- **pgAdmin**, a GUI maintained by members of the PostgreSQL community for developing and administering PostgreSQL databases on Linux, Unix, Mac OS X, and Windows.
- **Tora**, a free Oracle DBA interface with some PostgreSQL support.
- **DBeaver**, a multi-platform database tool that supports the most popular SQL and NoSQL databases, including Postgres. (Its Community Edition is free, while its Enterprise Edition is subject to a license fee.)
- **dbForge Studio for Postgres**, a user-friendly interface for creating, developing, and executing queries, editing and adjusting code, importing and exporting data, reporting and editing data, and more.
- **EMS SQL Manager for PostgreSQL**, with visual and text tools to build queries, a visual database designer, rapid database management, navigation and maintenance, and more.
- **JetBrains' DataGrip**, a multi-engine database IDE for writing SQL code, exploring databases, running queries, importing/exporting data, and more.

...it is unlikely that you will need to install and learn a new IDE, because chances are that the framework you are used to working with is already integrated with PostgreSQL.



EDB[™]
POSTGRES



...one of the design principles of PostgreSQL from the very start was that it would interact seamlessly and predictably with operating systems and storage frameworks.

EDB[™]
POSTGRES

2. Simple Local Setup and Configuration

Many traditional RDBMS interact in proprietary ways with storage and operating systems. For example, they require special kernels or raw devices, they use low-level operating system functionality, their backup tools and appliances are highly invasive, and so on. As a result, setting up the local environment and configuring proprietary database management systems are complex and time-consuming tasks that must be carried out by specialized and experienced DBAs before developers can get started on the project at hand.

By contrast, one of the design principles of PostgreSQL from the very start was that it would interact seamlessly and predictably with operating systems and storage frameworks. It is remarkably easy to install, set up, and configure a PostgreSQL deployment. In fact, developers can handle simpler deployments themselves, with DBAs only needing to get involved in larger-scale production deployments. PostgreSQL also integrates easily with standard operating system backup and monitoring tooling. It is lightweight and easy to move around, and surprises are rare. For all these reasons, PostgreSQL is a great enabler for developers across the entire application lifecycle—from prototyping to [production](#).

And if you would like to altogether avoid having to set up and configure a local environment for your PostgreSQL deployment, all of the major cloud providers offer PostgreSQL in a Database as a

Service (DBaaS) model, including [Amazon RDS](#) for PostgreSQL, [Azure Database for PostgreSQL](#), and [\(Google\) Cloud SQL for PostgreSQL](#). There are also third-party vendors that offer fully managed PostgreSQL cloud services, including our own [EDB Postgres Cloud Database Service](#) on the AWS cloud.

3. Flexible Multi-Model Architecture

The Postgres research project that ultimately spawned PostgreSQL was a pioneer of the object-relational DBMS (ORDBMS), i.e., a relational database that is extended to support some object-oriented programming features. These include extensibility of data types, access methods and functions, inheritance of table properties and data types, polymorphism (with an operator having different meanings within the same database), and encapsulation of tables. When storing and retrieving data in an ORDBMS, an object-relational mapping (ORM) tool automatically translates between non-scalar object values (attributes and fields) and scalar relational values (integers and strings organized within tables). Today, all of the leading ORM tools support PostgreSQL.

The multi-model ORDBMS approach is important for today's advanced apps that typically have to query and manipulate both structured and semi-structured data, which can require implementing both SQL and NoSQL databases in a single application. PostgreSQL, for example, has built-in support for storing JSON objects in a compact JSON binary format. This feature is especially

The Postgres research project that ultimately spawned PostgreSQL was a pioneer of the object-relational DBMS (ORDBMS)...

The logo for EnterpriseDB Postgres, featuring the letters "EDB" in a large, bold, white sans-serif font, with "POSTGRES" in a smaller, white sans-serif font directly below it. The background of the entire right side of the page is a dark blue, abstract image of a modern building's glass facade reflecting the sky and clouds.



Perhaps one of the biggest attractions of PostgreSQL for developers is that it lends itself well to dynamic DevOps environments and requirements.



convenient for applications that pass around semi-structured documents in JSON format and want to store them directly in the database in order to ensure high fidelity. This semi-structured data can be stored and queried alongside more structured data in tables as well, and all of it can be accessed using standard SQL and tooling.

In short, the power of the PostgreSQL multi-model architecture is that you do not have to embrace a new development and administrative tool chain to get the benefits of an ORDBMS—such as accelerated development cycles and enhanced runtime performance.

4. Agile and DevOps Friendly

Perhaps one of the biggest attractions of PostgreSQL for developers is that it lends itself well to dynamic DevOps environments and requirements.

First of all, getting started with PostgreSQL is frictionless. There are no lengthy procurement cycles, since the license for any kind of use, including commercial, is free and permissive. As already discussed above, setting up and configuring the PostgreSQL environment is straightforward. Plus, the fact that PostgreSQL is readily available as a cloud service aligns well with the cloud-based nature of DevOps methods.

Secondly, PostgreSQL's multi-model architecture and its support for JSON and ORM mean that

systems can be quickly prototyped without the need to fully design a detailed schema. In general, PostgreSQL fits well into continuous development, integration, and deployment processes, because it is easy to make and test small updates frequently. Developers can ensure consistent infrastructure configurations across the entire application lifecycle by introducing Infrastructure-as-Code recipes using Chef, Ansible, Puppet, or any other infrastructure automation platform. And PostgreSQL also fully supports the use of containers and microservices that are at the heart of today's modular, distributed app architectures.

Last but not least, PostgreSQL is highly portable, making it extremely easy to replicate instances for development, testing, and staging purposes. In PostgreSQL, there are clear boundaries between the database and underlying operating and storage systems. This is unlike traditional database management systems, such as DB2 and SQL Server that are inextricably tied into the DBMS owner's operating system. Thus, it is easy to move PostgreSQL databases around, run them locally or in the cloud, or run multiple versions that do not interfere with each other. Whatever you can do in your standard operating system or storage framework (copy directory, change permissions, etc.) you can do seamlessly in PostgreSQL. With everything, including data storage, installed in one directory and virtually no hooks to external resources, all you have to do is copy the file system in order to copy the database.

...it is easy to move
PostgreSQL
databases around,
run them locally or in
the cloud, or run
multiple versions that
do not interfere with
each other.

The logo for EnterpriseDB PostgreSQL. It features the letters "EDB" in a large, bold, white sans-serif font. Below "EDB" is the word "POSTGRES" in a smaller, white, all-caps sans-serif font. A small trademark symbol (TM) is located to the right of the "B" in "EDB". The background of the logo is a dark blue gradient with faint, glowing lines and a server rack visible in the lower right corner.

...developers who already have SQL skills will feel right at home in PostgreSQL.

5. Cool Things You Can Do With PostgreSQL

Before we get to the really cool stuff, we would like to point out that PostgreSQL is highly ANSI SQL-compliant. Thus, developers who already have SQL skills will feel right at home in PostgreSQL. Those who are new to SQL can take comfort in the fact that SQL is (relatively) easy to learn and that there's a lot of [PostgreSQL tooling](#) for building and testing SQL. EDB's [Postgres Enterprise Manager \(PEM\)](#), for example, extends the community's administration and management tool (pgAdmin) to meet the needs of large-scale enterprise PostgreSQL deployments, including a powerful SQL editor and query builder.

The really cool stuff starts, however, with the fact that, as discussed in the section on its multi-model architecture, PostgreSQL has built-in support for many non-relational NoSQL features without the drawbacks of NoSQL (such as client-side data analysis, lack of a powerful query language and optimizer, and data retention becoming an admin responsibility):

- **Arrays:** A table column can be defined as a multidimensional array of any built-in or user-defined base type, enum type, or composite type.
- **JSON/JSONB data type:** This validates stored values against the JSON rules, with built-in JSON-specific functions and operators available. JSON stores an exact copy of the input text, while JSONB decomposes the input text into a binary format that is faster to process and supports indexing.

EDB[™]
POSTGRES

- **Character Strings:** PostgreSQL has all the support that you would expect for storing, manipulating, and querying character strings. More advanced capabilities are available through extensions, such as full-text indexing using ZomboDB, an Elasticsearch PostgreSQL extension.
- **Full-Text Search:** This feature is fast and optimized, with support for fuzzy matching, ranking, phrase search, and multiple languages. It also supports GiST and GIN index types to speed up full-text searches.
- **Range Types:** Built-in range types include integer, bigint, numeric, timestamp (with and without time zone), and date. PostgreSQL also lets you define custom range types.
- **Geometric Support:** Data types for representing two-dimensional spatial objects as well as a set of functions and operators to perform geometric operations such as scaling, translation, rotation, and determining intersections.
- **XML Support:** The xml data type checks input values for well-formedness and provides functions to perform type-safe operations.
- **Many Specialized Index Types:** In addition to the GiST and GIN indexing infrastructures already mentioned in the context of full-text search, PostgreSQL supports:
 - **B-Tree:** Used for most data types and queries, can also be helpful in retrieving data in sorted order, and supports multicolumn indexes.
 - **SP-GiST:** Space-partitioned GiST, used for larger datasets with natural but uneven clustering.
 - **BRIN:** “Block Range Index,” a form of indexing for large datasets that line up sequentially.

The really cool stuff starts, however, with the fact that...

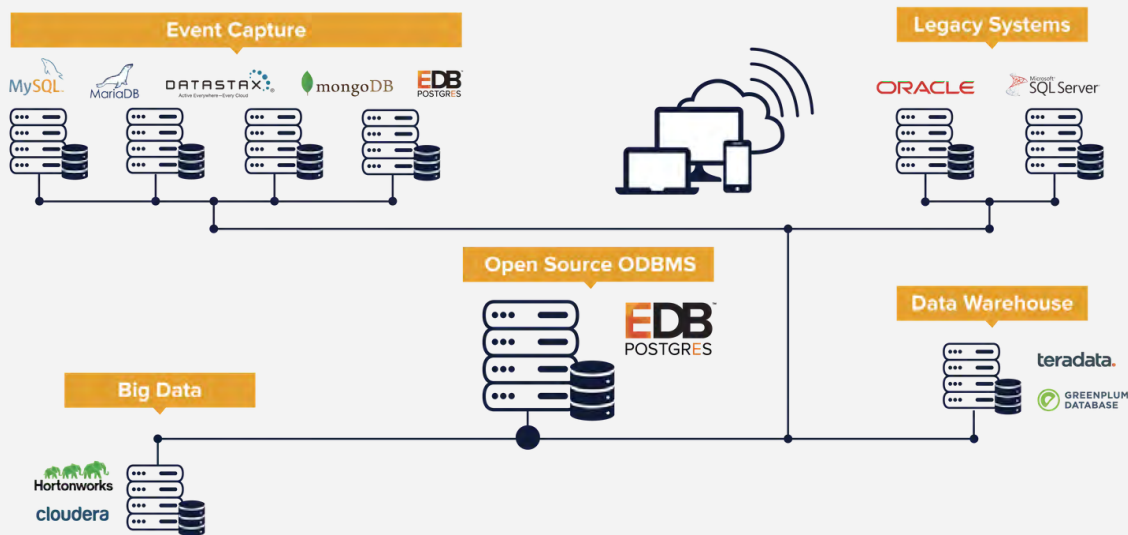
PostgreSQL has built-in support for many non-relational NoSQL features without the drawbacks of NoSQL



EDB[™]
POSTGRES

There are also numerous third-party plugins—many of them open source—that extend Postgres capabilities. These extensions are so tightly integrated that they behave the same way as built-in features. Some good open source examples are: [TimescaleDB](#) and [PipelineDB](#), for time-series data aggregation and graphs; [AgensGraph](#), a transactional graph database; and [PostGIS](#), which adds support for geographic objects and location queries, turning PostgreSQL into a fast, robust, and feature-rich spatial database management system.

PostgreSQL is a highly communicative and versatile DBMS that uses Foreign Data Wrappers (FDW) to seamlessly read from and write to foreign data sources. Some notable FDWs include CouchDB, Informix, MongoDB, MySQL, Neo4j, Oracle, and Redis. You can see a long list of currently available [FDWs in the Postgres Wiki](#) (although not all of them are officially supported by the PostgreSQL Global Development Group). With FDWs, you can use a single PostgreSQL database to federate at scale a diverse range of data sources and formats as well as app requirements, as illustrated below:



6. Smooth Road to Production

PostgreSQL's flexibility and ease of use also apply to how it is deployed in production. PostgreSQL lives comfortably on all of today's popular deployment platforms, from bare metal to VMs (on-prem or IaaS), containers, and DBaaS. The right choice of platform will depend on many factors, including the app's scalability and availability requirements, its architecture, and the extent to which the organization wants (or needs) to have direct control over infrastructure and orchestration issues.

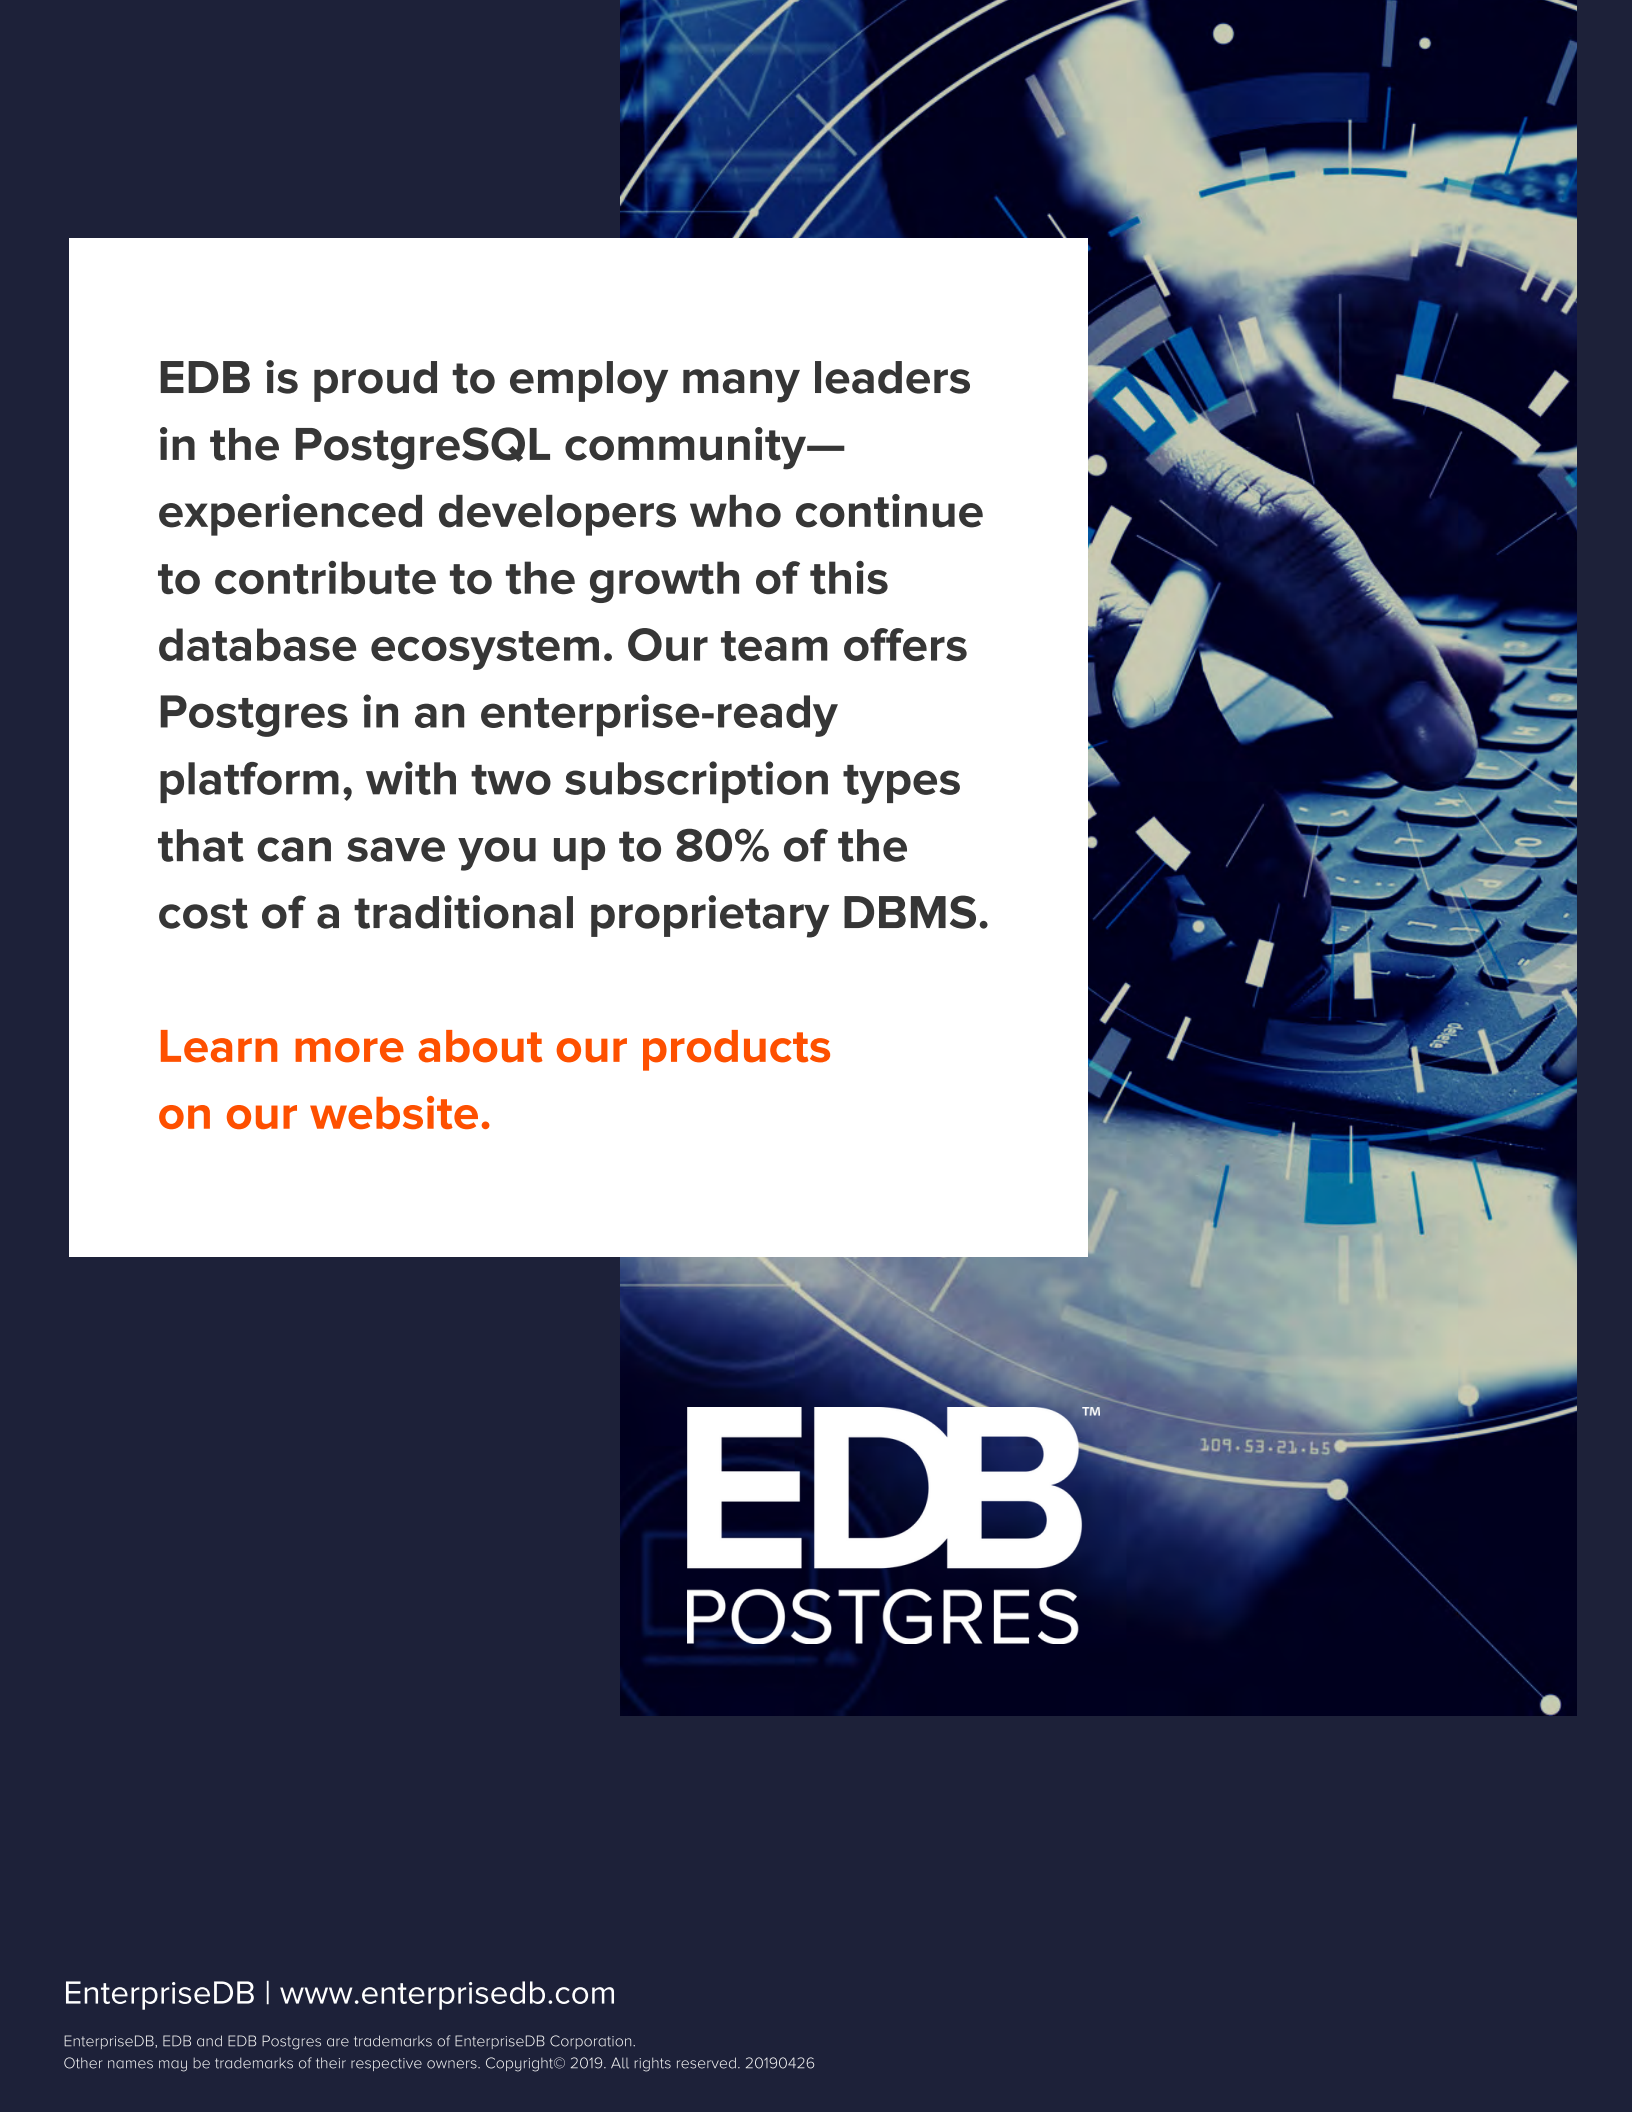
Today, for example, PostgreSQL is the third-most-popular technology being run by enterprises on Docker containers, [according to Datadog](#). Containerized PostgreSQL will behave consistently throughout the application lifecycle (development, staging, production) and is well-suited to microservice-based app architectures. It also delivers high availability in failover situations and optimizes compute-storage costs through elastic on-demand scaling rather than upfront over-provisioning.

Fully managed PostgreSQL DBaaS offerings, such as [EDB's Cloud Database Service](#) or other cloud service providers, simplify administration and can potentially boost developer productivity. A middle ground would be [EDB Postgres Ark Platform](#), which lets organizations set up their own DBaaS for flexible self-service database provisioning and also benefit from the enterprise-grade high availability, monitoring, backup, and load balancing of the EDB Postgres Advanced Server.

Final Note

Bruce Momjian is co-founder of the PostgreSQL Global Development Group, a tireless PostgreSQL evangelist, and a prolific blogger. He notes in his recent article [Making Postgres Central in Your Data Center](#) that all the reasons discussed above for developers to embrace PostgreSQL also perfectly position PostgreSQL to play a central role in the enterprise data center.

PostgreSQL's inherent reliability, flexibility, and extensibility, together with the support of its active open source community, ensure that it will always stay up to date as DBMS methods and technology stacks evolve.



EDB is proud to employ many leaders in the PostgreSQL community—experienced developers who continue to contribute to the growth of this database ecosystem. Our team offers Postgres in an enterprise-ready platform, with two subscription types that can save you up to 80% of the cost of a traditional proprietary DBMS.

Learn more about our products on our website.

EDBTM
POSTGRES