

Splendid Data white paper

PostgresPURE: 100% open source supported PostgreSQL, valuable for any type of organization

PostgresPURE highlights

Contents

Introduction	3
About PostgresPURE	
PostgresPURE	4
How is it composed?	
PostgresPURE components	
Benefits of PostgresPURE	
Release strategy and timing	
PostgresPURE support	
Contact	11
Ready to experience the benefits of PostgresPURE?	

Introduction

About PostgresPURE

PostgresPURE is cutting-edge software valuable for any type of organization using PostgreSQL (community version) as its core. By adding additional proven open source components to this, it offers you a very secure, highly available and future-proof 100% open source PostgreSQL database environment.

PostgreSQL is a powerful, open source object-relational database system with over 30 years of active development that has earned it a strong reputation for reliability, feature robustness, and performance.

In addition to being a very robust database, there's always a need for standard proven software components in a database environment for administration, backup/recovery, high availability, partition management and more. To make all these tasks effortless, we developed PostgresPURE.

PostgresPURE is 100% open source and contains besides PostgreSQL a set of open source components that makes it possible to set up a database environment that can compete with your Oracle database environments, or be even better. This means no more complicated license and maintenance contracts, no more restrictions regarding deployment and no vendor lock-in. Just 100% freedom.

PostgresPURE is fully equipped and is available on a subscription basis with 24-7 SLA technical support, in combination with a steady and predictable release policy.

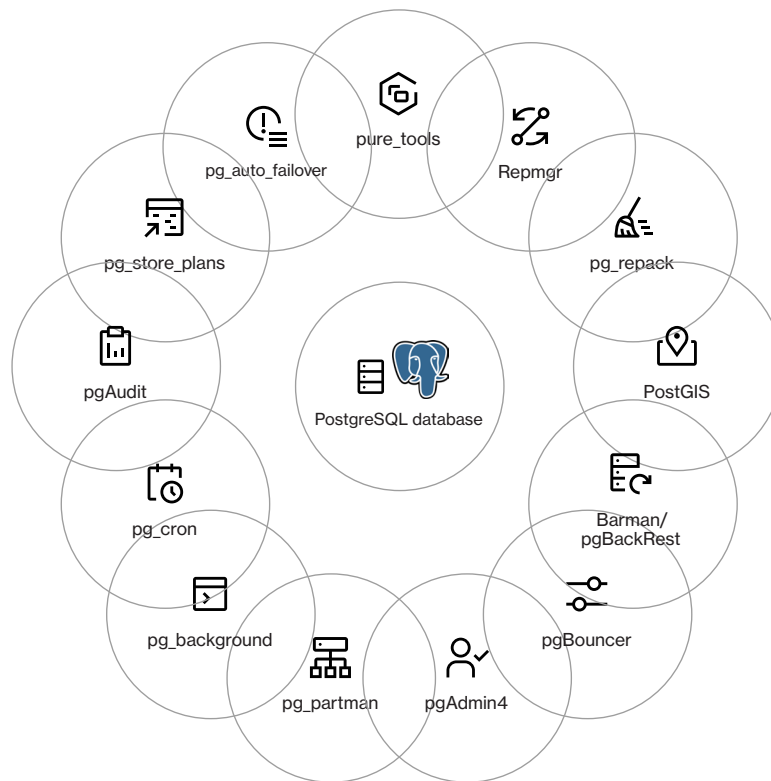


Some of our references that are convinced of what we promise and deliver

PostgresPURE

How is it composed?

PostgresPURE has been developed to provide you with as much functionality as possible on the basis of convenience and ease of use. This is then combined with premium support and release management.



PostgresPURE components

PostgresPURE is a subscription-based product. All components are 100% open source (free of charge) and in essence you only pay for the premium support and the release management. This is why the costs are just a fraction of what you are used to with your Oracle databases, where you have to pay one-off license fees and yearly maintenance fees.

Using our subscription model, you have access to our repo (via a certificate) in order to download the desired PostgresPURE release. In addition to the repo, in which the latest versions of PostgresPURE are available, there is also our archive. In this archive all the previous releases of PostgresPURE that we still support, are available to download.

Should you encounter any problems with PostgresPURE, we are available 24-7 through our online Splendid Data Help Center. You create a ticket and one of our experts will contact as soon as possible as defined in your SLA.

Support is a broad concept for us. We put our knowledge of PostgreSQL and all other components that are part of PostgresPURE at your disposal. It's why you also get access to the Splendid Data wiki. In this wiki you can find a lot of information to get the most out of PostgresPURE. Added to this, you can always ask our experts via the online Splendid Data Help Center for support. They are always on hand to answer all your queries.

PostgresPURE components



PostgreSQL database

The world's most advanced open source relational database.



pgAdmin4

A powerful database management tool built for ease of use.



pgBouncer

A leading connection pooler for setting up highly redundant environments.



pgAudit

A detailed session and/or object audit logging via the standard PostgreSQL facility.



Repmgr

A highly trusted replication and failover manager for PostgreSQL server clusters.



Barman/pgBackRest

Robust backup and recovery managers for planning backup and disaster recovery.



PostGIS

Provides spatial objects for storing information about location and mapping.



pg_background

Allows user to run arbitrary command in a background worker.



pg_repack

Removes bloat from tables and indexes to keep your database performing.



pg_partman

Creating and managing both time-based and serial-based table partition sets.



pure_tools

Shell scripts make your PostgreSQL life even easier.



pg_auto_failover

For monitoring and managing automated failover for a PostgreSQL cluster.



pg_cron

A simple job scheduler that enables you to add scheduled jobs to a database and execute SQL commands directly from a job.



pg_store_plans

For tracking execution plan statistics of all SQL statements executed by a server.

Benefits of PostgresPURE

- **Freedom** PostgresPURE is 100% open source to avoid any vendor lock-in. The licenses of all components are checked to prove they all are open source licensed and cannot cause any vendor lock-in.
- **Standardisation** PostgresPURE is setup to make it possible for you to maintain a standard for PostgreSQL and associated components, to prevent sprawling, and to allow you to stay in control.
- **Premium support** With a PostgresPURE subscription we guarantee 24-7 premium SLA-based technical support from our experts via the Splendid Data Help Center.
- **Release Management** Open source software means updates at the speed of innovation. Besides PostgreSQL, PostgresPURE consists of all kinds of components, each with their own release cycle. Splendid Data checks for every major and minor release of PostgreSQL to ensure these components remain in sync with PostgreSQL and the release policy of the PostgreSQL community.
- **Well tested** To provide you with a high quality product, PostgresPURE is thoroughly tested for both Intel x86_64 and IBM Power and for the different Linux versions. Splendid Data has test robots in place to check the configuration, the ability to install, the set-up (PostgreSQL in combination with each tool/ component and as a whole) and the functionalities.
- **Simplified installation** Your desired version of PostgresPURE can be downloaded from a central location (repo) via certified access. The scripts that take care of the installation of PostgresPURE create an unambiguous and consistent directory structure. Do you need an older version of PostgresPURE? This is taken care of through our PostgresPURE archive – where previous versions can be found.
- **Access to extensive knowledge** To get the most out of your PostgresPURE subscription you will get access to our extensive knowledge of PostgreSQL and all the components that are included in PostgresPURE. Access to this extensive knowledge is provided via our external wiki, reference documentation and the Splendid Data Help Center.
- **Speed of innovation** PostgreSQL becomes more and more a de-facto standard as a proven open source, object-relational database system. The speed of innovation accelerates within communities closely connected to PostgreSQL community. Very frequently, interesting new components become available that add yet more functionality to PostgreSQL and the ecosystem around PostgreSQL. At Splendid Data, we always keep an eye on these developments to assess if PostgresPURE should be extended to provide you with an even better product.
- **Flexibility** PostgresPURE can be deployed on bare metal or virtualised, and on premise or in the cloud (IaaS). All the flexibility you could need.

Release strategy and timing

The PostgreSQL community releases a major version of PostgreSQL once a year, usually at the end of September or beginning of October. Splendid Data starts preparing PostgresPURE as soon the first release of a new major (“x.0”) is available.

Around 6 weeks after a new major release, the first minor release (“x.1”) becomes available and also the new minor releases of the previous releases supported by the PostgreSQL community. Around 2 weeks after the first minor release of PostgreSQL we’ll have the new major release of PostgresPURE available. The new major release of PostgresPURE contains the first minor release of PostgreSQL. In a timeframe of around 6-8 weeks a new minor version of PostgresPURE is published with all additional components updated to the latest applicable version.

As you can see, PostgresPURE is in-sync with the regular release cycle of the PostgreSQL community. Every quarter we have the new minor versions of PostgresPURE available and if needed, new additional components.

Components	V6.22	V7.17	V8.13	V9.9	V10.4
PostgreSQL	11.21	12.16	13.12	14.9	15.4
Barman	3.6.0	3.6.0	3.6.0	3.6.0	3.6.0
Repmgr	5.4.1	5.4.1	5.4.1	5.4.1	5.4.1
pgBouncer	1.20.0	1.20.0	1.20.0	1.20.0	1.20.0
pgAudit	1.3.4	1.4.3	1.5.2	1.6.2	1.7.0
PostGIS	3.3.3	3.3.3	3.3.3	3.3.3	3.3.3
pgAdmin4 *	6.21	6.21	6.21	6.21	6.21
pgAdmin4 **	7.4	7.4	7.4	7.4	7.4
pgBackrest	2.45	2.45	2.45	2.45	2.45
pg_background	1.1	1.1	1.1	1.1	1.1
pg_repack	1.4.8	1.4.8	1.4.8	1.4.8	1.4.8
pg_partman	4.7.3	4.7.3	4.7.3	4.7.3	4.7.3
pure_tools	2.7	2.7	2.7	2.7	2.7
pg_store_plans	1.5	1.5	1.5	1.6.1	1.7.0
pg_auto_failover	2.0	2.0	2.0	2.0	2.0
pg_cron	1.5.2	1.5.2	1.5.2	1.5.2	1.5.2

* for Red Hat 7, CentOS 7 and SUSE 12

** for Red Hat 8, Rocky 8 and SUSE 15

PostgresPURE versions

Operating System	Architectures	PostgresPURE platforms					
		v6	v7	v8	v9	v10	
Red Hat - RHEL 7	X86_64 and ppc64le	✓	✓	✓	✓	✓	
Red Hat - RHEL 8	X86_64 and ppc64le	✓	✓	✓	✓	✓	
CentOS 7	X86_64 and ppc64le	✓	✓	✓	✓	✓	
CentOS 8	X86_64 and ppc64le	X	X	X	X	X	
SUSE - SLES 12	X86_64 and ppc64le	✓	✓	✓	✓	✓	
SUSE - SLES 15	X86_64 and ppc64le	✓	✓	✓	✓	✓	
Rocky Linux 8.7	X86_64	✓	✓	✓	✓	✓	

PostgresPURE platforms. Other Linux flavours; On request

PostgresPURE release support policy		
Current version	1st Release	Supported till
v6	October 2018	November 2023
v7	October 2019	November 2024
v8	October 2020	November 2025
v9	December 2021	November 2026

PostgresPURE release support policy

Support

Our Support Services consist of the following elements:

- **Installation support** Splendid Data provides you with the necessary information for installation of PostgresPURE.
- **Support documentation** Splendid Data provides you with the support documentation for PostgresPURE.
- **Splendid Data Help Center (SDHC)** Splendid Data makes its SDHC available to you and provides you with the necessary information on how to use it.
- **Incident management** In case of an incident with respect to PostgresPURE, you report the incident to Splendid Data, using the SDHC. You'll need to determine the incident level using SDHC in accordance with the following criteria:

Blocker levels

Blocker (level 1)

An incident that directly affects the continuation of work in a production environment, such as loss of production data and the non-functioning of the production environment. You cannot work with the application any longer due to the malfunctioning of the PostgreSQL database and/or one or more of the other components of PostgresPURE.

Critical (level 2)

An incident with the environment in which the operation of the PostgreSQL database and/or one or more of the other components of PostgresPURE in production environment is reduced severely and which has a large impact on (parts of) your business processes.

Major (level 3)

In the production environment: an incident that has no critical impact on your work in a production environment. You can continue working with the PostgreSQL database and the other components of PostgresPURE.

In the development environment: an incident that prevents you from working in your development environment any longer or migrating a tested PostgreSQL database into a production environment.

Minor (level 4)

In the production environment: error messages or warnings with no or negligible functional impact on your work in a production environment.

In the development environment: an incident that has no critical impact on your work in your development environment. You can continue working with the PostgreSQL database and the other components of PostgresPURE.

Trivial (level 5)

Reporting a documentation error or in textual description shown in the command line interface of the PostgreSQL database or any other components of PostgresPURE.

Service Levels

Splendid Data reacts and responds to any incident reported during the following service hours and within the following reaction and response times:

Service hours	Monday to Sunday, including public holidays in the country where End User is located, from 0.00 through 23.59	
Reaction and response times	Initial response within:	Ongoing response within:
Blocking (level 1)	1 hour, via SDHC	1 hour via SDHC or as otherwise agreed in writing
Critical (level 2)	2 hours, via SDHC	4 hours via SDHC or as otherwise agreed in writing
Major (level 3)	2 hours during Business Hours, via SDHC	4 hours during Business Hours via SDHC or as otherwise agreed in writing
Minor (level 4)	4 hours during Business Hours, via SDHC	8 hours during Business Hours via SDHC or as otherwise agreed in writing
Trivial (level 5)	4 hours during Business Hours, via SDHC	8 hours during Business Hours via SDHC or as otherwise agreed in writing

Splendid Data Service Levels

Generic questions

In case of a generic question to PostgresPURE, just send a request to us using the SDHC.

Request a change

In case of a change request (e.g. new functionality or adding a new component to PostgresPURE) with respect to PostgresPURE, use the SDHC to let us know.

Support Services for Development Purposes

Splendid Data provides you with advice on architecture, design, development and prototyping.

Contact

Ready to experience the benefits of PostgresPURE?

Reading about the benefits of PostgresPURE is useful, but there's nothing like a face-to-face meeting where we can explain everything in person and answer all your questions.

Please get in touch using the details below to arrange an initial meeting.

Michel Schöpgens

michel.schopgens@splendiddata.com

+31 85 773 19 99 (office)

+31 6 54 34 30 89 (mobile)

Pascal Boutin (for the French speaking markets)

pascal.boutin@splendiddata.com

+33 6 11 69 68 64 (mobile)

Splendid Data

Binnenhof 62a

1412 LC NAARDEN

The Netherlands

Find us online

www.splendiddata.com

www.linkedin.com/company/splendiddata

About Splendid Data

Splendid Data was founded in 2013 and since then our focus has been purely on PostgreSQL and to pass on its benefits to our clients and partners.

Our specialist team that has many years of in-depth knowledge of Oracle databases and PostgreSQL, enabling us to build **PostgresPURE** – our 100% open source subscription-based software that makes the most of PostgreSQL. At the same time, we developed **Cortex**, our unique solution for the automated migration of Oracle databases (DDL and DML) to native PostgreSQL. And to give organisations a full overview of their Oracle database environments and its suitability for migration, we created **Migration Scoping**, a clever piece of software that takes out all the guess work of a potential migration to PostgreSQL.

We're proud to assist companies of all shapes and sizes to find their way to digital liberty, agility and the power of PostgreSQL. Get in touch today to start your journey to a world without Oracle databases.

Confidential information owned by Splendid Data, to be used by the recipient only. This document, or any part of it, may not be reproduced, copied, circulated and/or distributed nor quoted without prior written approval from Splendid Data.

