

# Contents

- Introduction . . . . . 7
- Concepts . . . . . 7
  - Backup . . . . . 7
  - Restore . . . . . 8
  - Write Ahead Log (WAL) . . . . . 8
  - Encryption . . . . . 8
- Upgrading pgBackRest . . . . . 8
  - Upgrading pgBackRest from v1 to v2 . . . . . 8
- Build . . . . . 9
- Installation . . . . . 9
- Quick Start . . . . . 10
  - Setup Demo Cluster . . . . . 10
  - Configure Cluster Stanza . . . . . 11
  - Create the Repository . . . . . 12
  - Configure Archiving . . . . . 12
  - Configure Retention . . . . . 13
  - Configure Repository Encryption . . . . . 13
  - Create the Stanza . . . . . 13
  - Check the Configuration . . . . . 14
  - Perform a Backup . . . . . 14
  - Schedule a Backup . . . . . 15
  - Backup Information . . . . . 15
  - Restore a Backup . . . . . 16
- Backup . . . . . 16
  - Fast Start Option . . . . . 16
  - Archive Timeout . . . . . 17
- Monitoring . . . . . 17
  - In PostgreSQL . . . . . 20
  - Using jq . . . . . 21
- Retention . . . . . 22
  - Full Backup Retention . . . . . 22
  - Differential Backup Retention . . . . . 23
  - Archive Retention . . . . . 23
- Restore . . . . . 24
  - Delta Option . . . . . 25
  - Restore Selected Databases . . . . . 25
- Point-in-Time Recovery . . . . . 27
- S3 Support . . . . . 31
- Delete a Stanza . . . . . 33
- Dedicated Repository Host . . . . . 34
  - Installation . . . . . 34
  - Setup Passwordless SSH . . . . . 34

Configuration . . . . .	35
Perform a Backup . . . . .	36
Restore a Backup . . . . .	36
Parallel Backup / Restore . . . . .	36
Starting and Stopping . . . . .	37
Replication . . . . .	38
Installation . . . . .	38
Setup Passwordless SSH . . . . .	38
Hot Standby . . . . .	39
Streaming Replication . . . . .	41
Asynchronous Archiving . . . . .	43
Archive Push . . . . .	44
Archive Get . . . . .	45
Backup from a Standby . . . . .	46
Upgrading PostgreSQL . . . . .	47
Introduction . . . . .	50
Archive Get Command ( archive-get ) . . . . .	50
Command Options . . . . .	50
General Options . . . . .	50
Log Options . . . . .	53
Repository Options . . . . .	54
Stanza Options . . . . .	56
Archive Push Command ( archive-push ) . . . . .	56
Command Options . . . . .	57
General Options . . . . .	57
Log Options . . . . .	60
Repository Options . . . . .	61
Stanza Options . . . . .	63
Backup Command ( backup ) . . . . .	63
Command Options . . . . .	63
General Options . . . . .	66
Log Options . . . . .	68
Repository Options . . . . .	69
Stanza Options . . . . .	71
Check Command ( check ) . . . . .	73
Command Options . . . . .	73
General Options . . . . .	74
Log Options . . . . .	75
Repository Options . . . . .	77
Stanza Options . . . . .	79
Expire Command ( expire ) . . . . .	80
General Options . . . . .	80
Log Options . . . . .	82
Repository Options . . . . .	83

Stanza Options . . . . .	85
Help Command ( help ) . . . . .	86
Info Command ( info ) . . . . .	86
Command Options . . . . .	86
General Options . . . . .	86
Log Options . . . . .	88
Repository Options . . . . .	89
Restore Command ( restore ) . . . . .	91
Command Options . . . . .	91
General Options . . . . .	93
Log Options . . . . .	96
Repository Options . . . . .	97
Stanza Options . . . . .	99
Stanza Create Command ( stanza-create ) . . . . .	99
Command Options . . . . .	99
General Options . . . . .	100
Log Options . . . . .	102
Repository Options . . . . .	103
Stanza Options . . . . .	104
Stanza Delete Command ( stanza-delete ) . . . . .	106
Command Options . . . . .	106
General Options . . . . .	106
Log Options . . . . .	108
Repository Options . . . . .	109
Stanza Options . . . . .	110
Stanza Upgrade Command ( stanza-upgrade ) . . . . .	112
Command Options . . . . .	112
General Options . . . . .	112
Log Options . . . . .	114
Repository Options . . . . .	115
Stanza Options . . . . .	117
Start Command ( start ) . . . . .	118
General Options . . . . .	118
Log Options . . . . .	119
Repository Options . . . . .	121
Stanza Options . . . . .	123
Stop Command ( stop ) . . . . .	124
Command Options . . . . .	124
General Options . . . . .	124
Log Options . . . . .	125
Repository Options . . . . .	127
Stanza Options . . . . .	129
Version Command ( version ) . . . . .	130
Introduction . . . . .	130

Archive Options ( archive ) . . . . .	130
Asynchronous Archiving Option ( -archive-async ) . . . . .	130
Maximum Archive Get Queue Size Option ( -archive-get-queue-max ) . . . . .	130
Maximum Archive Push Queue Size Option ( -archive-push-queue-max ) . . . . .	131
Archive Timeout Option ( -archive-timeout ) . . . . .	131
Backup Options ( backup ) . . . . .	131
Check Archive Option ( -archive-check ) . . . . .	131
Copy Archive Option ( -archive-copy ) . . . . .	131
Backup from Standby Option ( -backup-standby ) . . . . .	132
Page Checksums Option ( -checksum-page ) . . . . .	132
Path/File Exclusions Option ( -exclude ) . . . . .	132
Manifest Save Threshold Option ( -manifest-save-threshold ) . . . . .	132
Resume Option ( -resume ) . . . . .	133
Start Fast Option ( -start-fast ) . . . . .	133
Stop Auto Option ( -stop-auto ) . . . . .	133
General Options ( general ) . . . . .	133
Buffer Size Option ( -buffer-size ) . . . . .	133
SSH client command Option ( -cmd-ssh ) . . . . .	133
Compress Option ( -compress ) . . . . .	134
Compress Level Option ( -compress-level ) . . . . .	134
Network Compress Level Option ( -compress-level-network ) . . . . .	134
Database Timeout Option ( -db-timeout ) . . . . .	134
Delta Option ( -delta ) . . . . .	134
Lock Path Option ( -lock-path ) . . . . .	134
Neutral Umask Option ( -neutral-umask ) . . . . .	135
Process Maximum Option ( -process-max ) . . . . .	135
Protocol Timeout Option ( -protocol-timeout ) . . . . .	135
Spool Path Option ( -spool-path ) . . . . .	135
Log Options ( log ) . . . . .	135
Console Log Level Option ( -log-level-console ) . . . . .	136
File Log Level Option ( -log-level-file ) . . . . .	136
Std Error Log Level Option ( -log-level-stderr ) . . . . .	136
Log Path Option ( -log-path ) . . . . .	137
Log Subprocesses Option ( -log-subprocess ) . . . . .	137
Log Timestamp Option ( -log-timestamp ) . . . . .	137
Repository Options ( repository ) . . . . .	137
Repository Cipher Passphrase Option ( -repo-cipher-pass ) . . . . .	137
Repository Cipher Type Option ( -repo-cipher-type ) . . . . .	137
Repository Hardlink Option ( -repo-hardlink ) . . . . .	138
Repository Host Option ( -repo-host ) . . . . .	138
Repository Host Command Option ( -repo-host-cmd ) . . . . .	138
Repository Host Configuration Option ( -repo-host-config ) . . . . .	138
Repository Host Configuration Include Path Option ( -repo-host-config-include-path ) . . . . .	138
Repository Host Configuration Path Option ( -repo-host-config-path ) . . . . .	138

Repository Host Port Option ( <code>-repo-host-port</code> ) . . . . .	139
Repository Host User Option ( <code>-repo-host-user</code> ) . . . . .	139
Repository Path Option ( <code>-repo-path</code> ) . . . . .	139
Archive Retention Option ( <code>-repo-retention-archive</code> ) . . . . .	139
Archive Retention Type Option ( <code>-repo-retention-archive-type</code> ) . . . . .	139
Differential Retention Option ( <code>-repo-retention-diff</code> ) . . . . .	140
Full Retention Option ( <code>-repo-retention-full</code> ) . . . . .	140
S3 Repository Bucket Option ( <code>-repo-s3-bucket</code> ) . . . . .	140
S3 SSL CA File Option ( <code>-repo-s3-ca-file</code> ) . . . . .	140
S3 SSL CA Path Option ( <code>-repo-s3-ca-path</code> ) . . . . .	140
S3 Repository Endpoint Option ( <code>-repo-s3-endpoint</code> ) . . . . .	140
S3 Repository Host Option ( <code>-repo-s3-host</code> ) . . . . .	141
S3 Repository Access Key Option ( <code>-repo-s3-key</code> ) . . . . .	141
S3 Repository Secret Access Key Option ( <code>-repo-s3-key-secret</code> ) . . . . .	141
S3 Repository Region Option ( <code>-repo-s3-region</code> ) . . . . .	141
S3 Repository Security Token Option ( <code>-repo-s3-token</code> ) . . . . .	141
S3 Repository Verify SSL Option ( <code>-repo-s3-verify-ssl</code> ) . . . . .	141
Repository Type Option ( <code>-repo-type</code> ) . . . . .	141
Restore Options ( <code>restore</code> ) . . . . .	142
Include Database Option ( <code>-db-include</code> ) . . . . .	142
Link All Option ( <code>-link-all</code> ) . . . . .	142
Link Map Option ( <code>-link-map</code> ) . . . . .	142
Recovery Option Option ( <code>-recovery-option</code> ) . . . . .	142
Tablespace Map Option ( <code>-tablespace-map</code> ) . . . . .	142
Map All Tablespaces Option ( <code>-tablespace-map-all</code> ) . . . . .	143
Stanza Options ( <code>stanza</code> ) . . . . .	143
PostgreSQL Host Option ( <code>-pg-host</code> ) . . . . .	143
PostgreSQL Host Command Option ( <code>-pg-host-cmd</code> ) . . . . .	143
PostgreSQL Host Configuration Option ( <code>-pg-host-config</code> ) . . . . .	143
PostgreSQL Host Configuration Include Path Option ( <code>-pg-host-config-include-path</code> ) . . . . .	143
PostgreSQL Host Configuration Path Option ( <code>-pg-host-config-path</code> ) . . . . .	144
PostgreSQL Host Port Option ( <code>-pg-host-port</code> ) . . . . .	144
PostgreSQL Host User Option ( <code>-pg-host-user</code> ) . . . . .	144
PostgreSQL Path Option ( <code>-pg-path</code> ) . . . . .	144
PostgreSQL Port Option ( <code>-pg-port</code> ) . . . . .	144
PostgreSQL Socket Path Option ( <code>-pg-socket-path</code> ) . . . . .	144
Introduction . . . . .	145
Current Stable Release . . . . .	145
v2.11 Release Notes . . . . .	145
Stable Releases . . . . .	145
v2.10 Release Notes . . . . .	145
v2.09 Release Notes . . . . .	146
v2.08 Release Notes . . . . .	146
v2.07 Release Notes . . . . .	147

v2.06 Release Notes	147
v2.05 Release Notes	148
v2.04 Release Notes	148
v2.03 Release Notes	149
v2.02 Release Notes	149
v2.01 Release Notes	150
v2.00 Release Notes	151
v1.29 Release Notes	151
v1.28 Release Notes	152
v1.27 Release Notes	153
v1.26 Release Notes	153
v1.25 Release Notes	154
v1.24 Release Notes	154
v1.23 Release Notes	155
v1.22 Release Notes	155
v1.21 Release Notes	155
v1.20 Release Notes	156
v1.19 Release Notes	156
v1.18 Release Notes	157
v1.17 Release Notes	157
v1.16 Release Notes	157
v1.15 Release Notes	157
v1.14 Release Notes	158
v1.13 Release Notes	158
v1.12 Release Notes	159
v1.11 Release Notes	160
v1.10 Release Notes	160
v1.09 Release Notes	160
v1.08 Release Notes	161
v1.07 Release Notes	161
v1.06 Release Notes	162
v1.05 Release Notes	162
v1.04 Release Notes	162
v1.03 Release Notes	163
v1.02 Release Notes	164
v1.01 Release Notes	164
v1.00 Release Notes	164
Pre-Stable Releases	165
v0.92 Release Notes	165
v0.91 Release Notes	165
v0.90 Release Notes	166
v0.89 Release Notes	166
v0.88 Release Notes	166
v0.87 Release Notes	167

v0.85 Release Notes	167
v0.82 Release Notes	167
v0.80 Release Notes	168
v0.78 Release Notes	169
v0.77 Release Notes	169
v0.75 Release Notes	169
v0.70 Release Notes	169
v0.65 Release Notes	170
v0.61 Release Notes	170
v0.60 Release Notes	171
v0.50 Release Notes	171
v0.30 Release Notes	171
v0.19 Release Notes	172
v0.18 Release Notes	172
v0.17 Release Notes	172
v0.16 Release Notes	172
v0.15 Release Notes	172
v0.11 Release Notes	172
v0.10 Release Notes	173

— title: User Guide - Debian & Ubuntu / PostgreSQL 10 draft: false —

## Introduction

This user guide is intended to be followed sequentially from beginning to end — each section depends on the last. For example, the Backup section relies on setup that is performed in the Quick Start section. Once pgBackRest is up and running then skipping around is possible but following the user guide in order is recommended the first time through.

Although the examples are targeted at Debian/Ubuntu and PostgreSQL 10, it should be fairly easy to apply this guide to any Unix distribution and PostgreSQL version. The only OS-specific commands are those to create, start, stop, and drop PostgreSQL clusters. The pgBackRest commands will be the same on any Unix system though the locations to install Perl libraries and executables may vary.

Configuration information and documentation for PostgreSQL can be found in the PostgreSQL [Manual](#) .

A somewhat novel approach is taken to documentation in this user guide. Each command is run on a virtual machine when the documentation is built from the XML source. This means you can have a high confidence that the commands work correctly in the order presented. Output is captured and displayed below the command when appropriate. If the output is not included it is because it was deemed not relevant or was considered a distraction from the narrative.

All commands are intended to be run as an unprivileged user that has sudo privileges for both the root and postgres users. It's also possible to run the commands directly as their respective users without modification and in that case the sudo commands can be stripped off.

## Concepts

The following concepts are defined as they are relevant to pgBackRest , PostgreSQL , and this user guide.

### Backup

A backup is a consistent copy of a database cluster that can be restored to recover from a hardware failure, to perform Point-In-Time Recovery, or to bring up a new standby.

**Full Backup** : pgBackRest copies the entire contents of the database cluster to the backup. The first backup of the database cluster is always a Full Backup. pgBackRest is always able to restore a full backup directly. The full backup does not depend on any files outside of the full backup for consistency.

**Differential Backup** : pgBackRest copies only those database cluster files that have changed since the last full backup. pgBackRest restores a differential backup by copying all of the files in the chosen differential backup and the appropriate unchanged files from the previous full backup. The advantage of a differential backup is that it requires less disk space than a full backup, however, the differential backup and the full backup must both be valid to restore the differential backup.

**Incremental Backup** : pgBackRest copies only those database cluster files that have changed since the last backup (which can be another incremental backup, a differential backup, or a full backup). As an incremental backup only includes those files changed since the prior backup, they are generally much smaller than full or differential backups. As with the differential backup, the incremental backup depends on other backups to be valid to restore the incremental backup. Since the incremental backup includes only those files since the last backup, all prior incremental backups back to the prior differential, the prior differential backup, and the prior full backup must all be valid to perform a restore of the incremental backup. If no differential backup exists then all prior incremental backups back to the prior full backup, which must exist, and the full backup itself must be valid to restore the incremental backup.

## Restore

A restore is the act of copying a backup to a system where it will be started as a live database cluster. A restore requires the backup files and one or more WAL segments in order to work correctly.

## Write Ahead Log (WAL)

WAL is the mechanism that PostgreSQL uses to ensure that no committed changes are lost. Transactions are written sequentially to the WAL and a transaction is considered to be committed when those writes are flushed to disk. Afterwards, a background process writes the changes into the main database cluster files (also known as the heap). In the event of a crash, the WAL is replayed to make the database consistent.

WAL is conceptually infinite but in practice is broken up into individual 16MB files called segments. WAL segments follow the naming convention 0000000100000A1E000000FE where the first 8 hexadecimal digits represent the timeline and the next 16 digits are the logical sequence number (LSN).

## Encryption

Encryption is the process of converting data into a format that is unrecognizable unless the appropriate password (also referred to as passphrase) is provided.

pgBackRest will encrypt the repository based on a user-provided password, thereby preventing unauthorized access to data stored within the repository.

## Upgrading pgBackRest

### Upgrading pgBackRest from v1 to v2

Upgrading from v1 to v2 is fairly straight-forward. The repository format has not changed and all non-deprecated options from v1 are accepted, so for most installations it is simply a matter of installing the new version.

However, there are a few caveats:

- The deprecated thread-max option is no longer valid. Use process-max instead.
- The deprecated archive-max-mb option is no longer valid. This has been replaced with the archive-push-queue-max option which has different semantics.
- The default for the backup-user option has changed from backrest to pgbackrest .
- In v2.02 the default location of the pgBackRest configuration file has changed from /etc/pgbackrest.conf to /etc/pgbackrest/pgbackrest.conf . If /etc/pgbackrest/pgbackrest.conf does not exist, the /etc/pgbackrest.conf file will be loaded instead, if it exists.

Many option names have changed to improve consistency although the old names from v1 are still accepted. In general, db-\* options have been renamed to pg-\* and backup-\* / retention-\* options have been renamed to repo-\* when appropriate.

PostgreSQL and repository options must be indexed when using the new names introduced in v2 , e.g. pg1-host , pg1-path , repo1-path , repo1-type , etc. Only one repository is allowed currently but more flexibility is planned for v2 .



## Build

Debian/Ubuntu packages for pgBackRest are available at [apt.postgresql.org](http://apt.postgresql.org) . If they are not provided for your distribution/version it is easy to download the source and install manually.

When building from source it is best to use a build host rather than building on production. Many of the tools required for the build should generally not be installed in production. pgBackRest consists of a single executable so it is easy to copy to a new host once it is built.

build Download version 2.11 of pgBackRest

```
sudo wget -q -O - \
  https://github.com/pgbackrest/pgbackrest/archive/release/2.11.tar.gz | \
  sudo tar zx -C /root
```

build Install build dependencies

```
sudo apt-get install build-essential libssl-dev libxml2-dev libperl-dev zlib1g-dev
```

pgBackRest supports 32-bit distributions that build Perl with 64-bit integer support.

build Check for 64-bit integers

```
perl -V | grep USE_64_BIT_INT
```

```
USE_64_BIT_ALL USE_64_BIT_INT USE_ITHREADS
```

The pgBackRest executable is written in C. This allows certain time-critical commands (like `async archive-push / archive-get` ) to run more quickly.

build Build pgBackRest package

```
sudo make -s -C /root/pgbackrest-release-2.11/src
```

## Installation

A new host named `pg1` is created to contain the demo cluster and run pgBackRest examples.

pgBackRest needs to be installed from a package or installed manually as shown here.

pg-primary Copy pgBackRest binary from build host

```
sudo scp build:/root/pgbackrest-release-2.11/src/pgbackrest /usr/bin
```

```
sudo chmod 755 /usr/bin/pgbackrest
```

pgBackRest contains embedded Perl which requires some additional modules.

pg-primary Install required Perl packages

```
sudo apt-get install libdbd-pg-perl
```

Finally, pgBackRest requires log and configuration directories and a configuration file.

pg-primary Create pgBackRest configuration file and directories

```
sudo mkdir -p -m 770 /var/log/pgbackrest
```

```
sudo chown postgres:postgres /var/log/pgbackrest
```

```
sudo mkdir -p /etc/pgbackrest
```

```
sudo mkdir -p /etc/pgbackrest/conf.d
```

```
sudo touch /etc/pgbackrest/pgbackrest.conf
```

```
sudo chmod 640 /etc/pgbackrest/pgbackrest.conf
```

```
sudo chown postgres:postgres /etc/pgbackrest/pgbackrest.conf
```

pgBackRest should now be properly installed but it is best to check. If any dependencies were missed then you will get an error when running pgBackRest from the command line.

pg-primary Make sure the installation worked

```
sudo -u postgres pgbackrest
```

```
pgBackRest 2.11 - General help
```

Usage:

```
pgbackrest [options] [command]
```

Commands:

```
archive-get      Get a WAL segment from the archive.
archive-push     Push a WAL segment to the archive.
backup           Backup a database cluster.
check            Check the configuration.
expire           Expire backups that exceed retention.
help             Get help.
info             Retrieve information about backups.
restore          Restore a database cluster.
stanza-create    Create the required stanza data.
stanza-delete    Delete a stanza.
stanza-upgrade   Upgrade a stanza.
start            Allow pgBackRest processes to run.
stop             Stop pgBackRest processes from running.
version          Get version.
```

Use 'pgbackrest help [command]' for more information.

## Quick Start

The Quick Start section will cover basic configuration of pgBackRest and PostgreSQL and introduce the backup , restore , and info commands.

## Setup Demo Cluster

Creating the demo cluster is optional but is strongly recommended, especially for new users, since the example commands in the user guide reference the demo cluster; the examples assume the demo cluster is running on the default port (i.e. 5432). The cluster will not be started until a later section because there is still some configuration to do.

pg-primary Create the demo cluster

```
sudo -u postgres /usr/lib/postgresql/10/bin/initdb \
-D /var/lib/postgresql/10/demo -k -A peer
```

```
sudo pg_createcluster 10 demo
```

```
Configuring already existing cluster (configuration: /etc/postgresql/10/demo, data:
/var/lib/postgresql/10/demo, owner: 106:110)
```

Ver	Cluster	Port	Status	Owner	Data directory	Log file
10	demo	5432	down	postgres	/var/lib/postgresql/10/demo	/var/log/postgresql/postgresql-10-demo.log

By default PostgreSQL will only accept local connections. The examples in this guide will require connections from other servers so listen\_addresses is configured to listen on all interfaces. This may not be appropriate for secure installations.

pg-primary : /etc/postgresql/10/demo/postgresql.conf Set listen\_addresses

```
listen_addresses = '*'
```

For demonstration purposes the log\_line\_prefix setting will be minimally configured. This keeps the log output as brief as possible to better illustrate important information.

pg-primary : /etc/postgresql/10/demo/postgresql.conf Set log\_line\_prefix

```
listen_addresses = '*'
```

```
log_line_prefix = "
```

## Configure Cluster Stanza

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as app or dw, rather than the local cluster name, such as main or prod.

The name 'demo' describes the purpose of this cluster accurately so that will also make a good stanza name.

pgBackRest needs to know where the base data directory for the PostgreSQL cluster is located. The path can be requested from PostgreSQL directly but in a recovery scenario the PostgreSQL process will not be available. During backups the value supplied to pgBackRest will be compared against the path that PostgreSQL is running on and they must be equal or the backup will return an error. Make sure that pg-path is exactly equal to data\_directory in postgresql.conf .

By default Debian/Ubuntu stores clusters in /var/lib/postgresql/[version]/[cluster] so it is easy to determine the correct path for the data directory.

When creating the /etc/pgbackrest/pgbackrest.conf file, the database owner (usually postgres ) must be granted read privileges.

pg-primary : /etc/pgbackrest/pgbackrest.conf Configure the PostgreSQL cluster data directory

```
[demo]
pg1-path=/var/lib/postgresql/10/demo
```

pgBackRest configuration files follow the Windows INI convention. Sections are denoted by text in brackets and key/value pairs are contained in each section. Lines beginning with # are ignored and can be used as comments.

There are multiple ways the pgBackRest configuration files can be loaded:

- config and config-include-path are default: the default config file will be loaded, if it exists, and \*.conf files in the default config include path will be appended, if they exist.
- config option is specified: only the specified config file will be loaded and is expected to exist.
- config-include-path is specified: \*.conf files in the config include path will be loaded and the path is required to exist. The default config file will be loaded if it exists. If it is desirable to load only the files in the specified config include path, then the -no-config option can also be passed.
- config and config-include-path are specified: using the user-specified values, the config file will be loaded and \*.conf files in the config include path will be appended. The files are expected to exist.
- config-path is specified: this setting will override the base path for the default location of the config file and/or the base path of the default config-include-path setting unless the config and/or config-include-path option is explicitly set.

The files are concatenated as if they were one big file; order doesn't matter, but there is precedence based on sections. The precedence (highest to lowest) is:

- [ stanza : command ]
- [ stanza ]
- [global: command ]
- [global]

NOTE:

-config , -config-include-path and -config-path are command-line only options.

pgBackRest can also be configured using environment variables as described in the [command reference](#) .

pg-primary Configure log-path using the environment

```
sudo -u postgres bash -c ' \
    export PGBACKREST_LOG_PATH=/path/set/by/env && \
    pgbackrest --log-level-console=error help backup log-path'
```

```
pgBackRest 2.11 - 'backup' command - 'log-path' option help
```

Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if log-level-file=off then no log path is required.

```
current: /path/set/by/env
```

```
default: /var/log/pgbackrest
```

## Create the Repository

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

For this demonstration the repository will be stored on the same host as the PostgreSQL server. This is the simplest configuration and is useful in cases where traditional backup software is employed to backup the database host.

pg-primary Create the pgBackRest repository

```
sudo mkdir -p /var/lib/pgbackrest
```

```
sudo chmod 750 /var/lib/pgbackrest
```

```
sudo chown postgres:postgres /var/lib/pgbackrest
```

The repository path must be configured so pgBackRest knows where to find it.

pg-primary : /etc/pgbackrest/pgbackrest.conf Configure the pgBackRest repository path

```
[demo]
pg1-path=/var/lib/postgresql/10/demo
```

```
[global]
repo1-path=/var/lib/pgbackrest
```

## Configure Archiving

Backing up a running PostgreSQL cluster requires WAL archiving to be enabled. Note that *at least* one WAL segment will be created during the backup process even if no explicit writes are made to the cluster.

pg-primary : /etc/postgresql/10/demo/postgresql.conf Configure archive settings

```
archive_command = 'pgbackrest -stanza=demo archive-push %p'
archive_mode = on
listen_addresses = '*'
log_line_prefix = ''
max_wal_senders = 3
wal_level = hot_standby
```

The wal\_level setting must be set to archive at a minimum but hot\_standby and logical also work fine for backups. Setting wal\_level to hot\_standby and increasing max\_wal\_senders is a good idea even if you do not currently run a hot standby as this will allow them to be added later without restarting the primary cluster.

The PostgreSQL cluster must be restarted after making these changes and before performing a backup.

pg-primary Restart the demo cluster

```
sudo pg_ctlcluster 10 demo restart
```

When archiving a WAL segment is expected to take more than 60 seconds (the default) to reach the pgBackRest repository, then the pgBackRest archive-timeout option should be increased. Note that this option is not the same as the PostgreSQL archive\_timeout option which is used to force a WAL segment switch; useful for databases where there are long periods of inactivity. For more information on the PostgreSQL archive\_timeout option, see PostgreSQL [Write Ahead Log](#).

The archive-push command can be configured with its own options. For example, a lower compression level may be set to speed archiving without affecting the compression used for backups.

pg-primary : /etc/pgbackrest/pgbackrest.conf Config archive-push to use a lower compression level

```
[demo]
pg1-path=/var/lib/postgresql/10/demo
```

```
[global]
repo1-path=/var/lib/pgbackrest
```

```
[global:archive-push]
compress-level=3
```

This configuration technique can be used for any command and can even target a specific stanza, e.g. `demo:archive-push` .

## Configure Retention

`pgBackRest` expires backups based on retention options.

```
pg-primary : /etc/pgbackrest/pgbackrest.conf  Configure retention to 2 full backups
```

```
[demo]
pg1-path=/var/lib/postgresql/10/demo
```

```
[global]
repo1-path=/var/lib/pgbackrest
repo1-retention-full=2
```

```
[global:archive-push]
compress-level=3
```

More information about retention can be found in the Retention section.

## Configure Repository Encryption

The repository will be configured with a cipher type and key to demonstrate encryption. Encryption is always performed client-side even if the repository type (e.g. `S3`) supports encryption.

It is important to use a long, random passphrase for the cipher key. A good way to generate one is to run: `openssl rand -base64 48` .

```
pg-primary : /etc/pgbackrest/pgbackrest.conf  Configure pgBackRest repository encryption
```

```
[demo]
pg1-path=/var/lib/postgresql/10/demo
```

```
[global]
repo1-cipher-pass=zWaf6XtpjIVZC5444yXB+cgFDF17MxGlGkZSaoPvTGirhPygu4jOKOXf9LO4vjfO
repo1-cipher-type=aes-256-cbc
repo1-path=/var/lib/pgbackrest
repo1-retention-full=2
```

```
[global:archive-push]
compress-level=3
```

Once the repository has been configured and the stanza created and checked, the repository encryption settings cannot be changed.

## Create the Stanza

The `stanza-create` command must be run on the host where the repository is located to initialize the stanza. It is recommended that the `check` command be run after `stanza-create` to ensure archiving and backups are properly configured.

```
pg-primary  Create the stanza and check the configuration
```

```
sudo -u postgres pgbackrest --stanza=demo --log-level-console=info stanza-create
```

```
P00  INFO: stanza-create command begin 2.11: --log-level-console=info --log-level-stderr=off
--no-log-timestamp --pg1-path=/var/lib/postgresql/10/demo --repo1-cipher-pass=
--repo1-cipher-type=aes-256-cbc --repo1-path=/var/lib/pgbackrest --stanza=demo
```

```
P00  INFO: stanza-create command end: completed successfully
```

## Check the Configuration

The check command validates that pgBackRest and the archive\_command setting are configured correctly for archiving and backups. It detects misconfigurations, particularly in archiving, that result in incomplete backups because required WAL segments did not reach the archive. The command can be run on the database or the repository host. The command may also be run on the standby host, however, since pg\_switch\_xlog() / pg\_switch\_wal() cannot be performed on the standby, the command will only test the repository configuration.

Note that pg\_create\_restore\_point('pgBackRest Archive Check') and pg\_switch\_xlog() / pg\_switch\_wal() are called to force PostgreSQL to archive a WAL segment. Restore points are only supported in PostgreSQL >= 9.1 so for older versions the check command may fail if there has been no write activity since the last log rotation, therefore it is recommended that activity be generated by the user if there have been no writes since the last WAL switch before running the check command.

pg-primary Check the configuration

```
sudo -u postgres pgbackrest --stanza=demo --log-level-console=info check
```

```
P00 INFO: check command begin 2.11: --log-level-console=info --log-level-stderr=off
--no-log-timestamp --pg1-path=/var/lib/postgresql/10/demo --repo1-cipher-pass=
--repo1-cipher-type=aes-256-cbc --repo1-path=/var/lib/pgbackrest --stanza=demo
```

```
P00 INFO: WAL segment 000000010000000000000001 successfully stored in the archive at
'/var/lib/pgbackrest/archive/demo/10-1/0000000100000000/000000010000000000000001-f11824666a888d45a
```

```
P00 INFO: check command end: completed successfully
```

## Perform a Backup

To perform a backup of the PostgreSQL cluster run pgBackRest with the backup command.

pg-primary Backup the demo cluster

```
sudo -u postgres pgbackrest --stanza=demo \
--log-level-console=info backup
```

```
P00 INFO: backup command begin 2.11: --log-level-console=info --log-level-stderr=off
--no-log-timestamp --pg1-path=/var/lib/postgresql/10/demo --repo1-cipher-pass=
--repo1-cipher-type=aes-256-cbc --repo1-path=/var/lib/pgbackrest --repo1-retention-full=2
--stanza=demo
```

```
P00 WARN: no prior backup exists, incr backup has been changed to full
```

```
P00 INFO: execute non-exclusive pg_start_backup() with label "pgBackRest backup started at
2019-03-10 12:11:48": backup begins after the next regular checkpoint completes
```

```
P00 INFO: backup start archive = 000000010000000000000002, lsn = 0/2000028
[filtered 941 lines of output]
```

```
P01 INFO: backup file /var/lib/postgresql/10/demo/base/1/12820 (0B, 100%)
```

```
P01 INFO: backup file /var/lib/postgresql/10/demo/base/1/12815 (0B, 100%)
```

```
P00 INFO: full backup size = 22.4MB
```

```
P00 INFO: execute non-exclusive pg_stop_backup() and wait for all WAL segments to archive
```

```
P00 INFO: backup stop archive = 000000010000000000000002, lsn = 0/2000130
[filtered 4 lines of output]
```

By default pgBackRest will attempt to perform an incremental backup. However, an incremental backup must be based on a full backup and since no full backup existed pgBackRest ran a full backup instead.

The type option can be used to specify a full or differential backup.

pg-primary Differential backup of the demo cluster

```
sudo -u postgres pgbackrest --stanza=demo --type=diff \
--log-level-console=info backup
```

```
[filtered 4 lines of output]
```

```
P01 INFO: backup file /var/lib/postgresql/10/demo/global/pg_control (8KB, 99%) checksum
40b1b49d010202c2752f6a39eb72e695d0025077
```

```
P01 INFO: backup file /var/lib/postgresql/10/demo/pg_logical/replorigin_checkpoint (8B, 100%)
checksum 347fc8f2df71bd4436e38bd1516ccd7ea0d46532
```

```
P00 INFO: diff backup size = 8KB
```

```
P00 INFO: execute non-exclusive pg_stop_backup() and wait for all WAL segments to archive
P00 INFO: backup stop archive = 00000001000000000000000003, lsn = 0/30000F8
      [filtered 4 lines of output]
```

This time there was no warning because a full backup already existed. While incremental backups can be based on a full *or* differential backup, differential backups must be based on a full backup. A full backup can be performed by running the backup command with `-type=full`.

More information about the backup command can be found in the Backup section.

## Schedule a Backup

Backups can be scheduled with utilities such as cron.

In the following example, two cron jobs are configured to run; full backups are scheduled for 6:30 AM every Sunday with differential backups scheduled for 6:30 AM Monday through Saturday. If this crontab is installed for the first time mid-week, then pgBackRest will run a full backup the first time the differential job is executed, followed the next day by a differential backup.

```
#m h dom mon dow command
30 06 * * 0 pgbackrest --type=full --stanza=demo backup
30 06 * * 1-6 pgbackrest --type=diff --stanza=demo backup
```

Once backups are scheduled it's important to configure retention so backups are expired on a regular schedule, see Retention .

## Backup Information

Use the info command to get information about backups.

```
pg-primary Get info for the demo cluster
```

```
sudo -u postgres pgbackrest info
```

```
stanza: demo
  status: ok
  cipher: aes-256-cbc

  db (current)
    wal archive min/max (10-1): 000000010000000000000002/000000010000000000000003
```

```
  full backup: 20190310-121148F
```

```
    timestamp start/stop: 2019-03-10 12:11:48 / 2019-03-10 12:11:55
    wal start/stop: 000000010000000000000002 / 000000010000000000000002
    database size: 22.4MB, backup size: 22.4MB
    repository size: 2.7MB, repository backup size: 2.7MB
```

```
  diff backup: 20190310-121148F_20190310-121156D
```

```
    timestamp start/stop: 2019-03-10 12:11:56 / 2019-03-10 12:12:00
    wal start/stop: 000000010000000000000003 / 000000010000000000000003
    database size: 22.4MB, backup size: 8.2KB
    repository size: 2.7MB, repository backup size: 517B
    backup reference list: 20190310-121148F
```

Each stanza has a separate section and it is possible to limit output to a single stanza with the `--stanza` option. The stanza ' status ' gives a brief indication of the stanza's health. If this is ' ok ' then pgBackRest is functioning normally. The ' wal archive min/max ' shows the minimum and maximum WAL currently stored in the archive. Note that there may be gaps due to archive retention policies or other reasons.

The backups are displayed oldest to newest. The oldest backup will *always* be a full backup (indicated by an F at the end of the label) but the newest backup can be full, differential (ends with D ), or incremental (ends with I ).

The ' timestamp start/stop ' defines the time period when the backup ran. The ' timestamp stop ' can be used to determine the backup to use when performing Point-In-Time Recovery. More information about Point-In-Time Recovery can be found in the Point-In-Time Recovery section.

The ' wal start/stop ' defines the WAL range that is required to make the database consistent when restoring. The backup command will ensure that this WAL range is in the archive before completing.

The ' database size ' is the full uncompressed size of the database while ' backup size ' is the amount of data actually backed up (these will be the same for full backups). The ' repository size ' includes all the files from this backup and any referenced backups that are required to restore the database while ' repository backup size ' includes only the files in this backup (these will also be the same for full backups). Repository sizes reflect compressed file sizes if compression is enabled in pgBackRest or the filesystem.

The ' backup reference list ' contains the additional backups that are required to restore this backup.

## Restore a Backup

Backups can protect you from a number of disaster scenarios, the most common of which are hardware failure and data corruption. The easiest way to simulate data corruption is to remove an important PostgreSQL cluster file.

pg-primary Stop the demo cluster and delete the pg\_control file

```
sudo pg_ctlcluster 10 demo stop
```

```
sudo -u postgres rm /var/lib/postgresql/10/demo/global/pg_control
```

Starting the cluster without this important file will result in an error.

pg-primary Attempt to start the corrupted demo cluster

```
sudo pg_ctlcluster 10 demo start
```

```
Error: /usr/lib/postgresql/10/bin/pg_ctl /usr/lib/postgresql/10/bin/pg_ctl start -D
/var/lib/postgresql/10/demo -l /var/log/postgresql/postgresql-10-demo.log -s -o -c
config_file="/etc/postgresql/10/demo/postgresql.conf" exited with status 1:
```

```
postgres: could not find the database system
```

```
Expected to find it in the directory "/var/lib/postgresql/10/demo",
but could not open file "/var/lib/postgresql/10/demo/global/pg_control": No such file or directory
Examine the log output.
```

To restore a backup of the PostgreSQL cluster run pgBackRest with the restore command. The cluster needs to be stopped (in this case it is already stopped) and all files must be removed from the PostgreSQL data directory.

pg-primary Remove old files from demo cluster

```
sudo -u postgres find /var/lib/postgresql/10/demo -mindepth 1 -delete
```

pg-primary Restore the demo cluster and start PostgreSQL

```
sudo -u postgres pgbackrest --stanza=demo restore
```

```
sudo pg_ctlcluster 10 demo start
```

This time the cluster started successfully since the restore replaced the missing pg\_control file.

More information about the restore command can be found in the Restore section.

## Backup

The Backup section introduces additional backup command features.

### Fast Start Option

By default pgBackRest will wait for the next regularly scheduled checkpoint before starting a backup. Depending on the checkpoint\_timeout and checkpoint\_segments settings in PostgreSQL it may be quite some time before a checkpoint completes and the backup can begin.

pg-primary Incremental backup of the demo cluster with the regularly scheduled checkpoint

```
sudo -u postgres pgbackrest --stanza=demo --type=incr \
--log-level-console=info backup
```



```
P00 INFO: backup command begin 2.11: --log-level-console=info --log-level-stderr=off
--no-log-timestamp --pg1-path=/var/lib/postgresql/10/demo --repo1-cipher-pass=
--repo1-cipher-type=aes-256-cbc --repo1-path=/var/lib/pgbackrest --repo1-retention-full=2
--stanza=demo --type=incr
P00 INFO: last backup label = 20190310-121148F_20190310-121156D, version = 2.11

P00 INFO: execute non-exclusive pg_start_backup() with label "pgBackRest backup started at
2019-03-10 12:12:11": backup begins after the next regular checkpoint completes

P00 INFO: backup start archive = 00000002000000000000000005, lsn = 0/5000028
P00 WARN: a timeline switch has occurred since the last backup, enabling delta checksum
[filtered 8 lines of output]
```

When `-start-fast` is passed on the command-line or `start-fast=y` is set in `/etc/pgbackrest/pgbackrest.conf` an immediate checkpoint is requested and the backup will start more quickly. This is convenient for testing and for ad-hoc backups. For instance, if a backup is being taken at the beginning of a release window it makes no sense to wait for a checkpoint. Since regularly scheduled backups generally only happen once per day it is unlikely that enabling the `start-fast` in `/etc/pgbackrest/pgbackrest.conf` will negatively affect performance, however for high-volume transactional systems you may want to pass `-start-fast` on the command-line instead. Alternately, it is possible to override the setting in the configuration file by passing `-no-start-fast` on the command-line.

`pg-primary` : `/etc/pgbackrest/pgbackrest.conf` Enable the `start-fast` option

```
[demo]
pg1-path=/var/lib/postgresql/10/demo

[global]
repo1-cipher-pass=zWaf6XtpjIVZC5444yXB+cgFDF17MxGlgkZSaoPvTGirhPygu4jOKOXf9LO4vjfO
repo1-cipher-type=aes-256-cbc
repo1-path=/var/lib/pgbackrest
repo1-retention-full=2
start-fast=y
```

```
[global:archive-push]
compress-level=3
```

`pg-primary` Incremental backup of the demo cluster with an immediate checkpoint

```
sudo -u postgres pgbackrest --stanza=demo --type=incr \
--log-level-console=info backup
```

```
P00 INFO: backup command begin 2.11: --log-level-console=info --log-level-stderr=off
--no-log-timestamp --pg1-path=/var/lib/postgresql/10/demo --repo1-cipher-pass=
--repo1-cipher-type=aes-256-cbc --repo1-path=/var/lib/pgbackrest --repo1-retention-full=2
--stanza=demo --start-fast --type=incr
P00 INFO: last backup label = 20190310-121148F_20190310-121211I, version = 2.11

P00 INFO: execute non-exclusive pg_start_backup() with label "pgBackRest backup started at
2019-03-10 12:12:18": backup begins after the requested immediate checkpoint completes

P00 INFO: backup start archive = 00000002000000000000000006, lsn = 0/6000028
P01 INFO: backup file /var/lib/postgresql/10/demo/global/pg_control (8KB, 99%) checksum
b77df9e069b4e1854a865fdf0631e443edf09e3a
[filtered 8 lines of output]
```

## Archive Timeout

During an online backup `pgBackRest` waits for WAL segments that are required for backup consistency to be archived. This wait time is governed by the `pgBackRest` `archive-timeout` option which defaults to 60 seconds. If archiving an individual segment is known to take longer then this option should be increased.

## Monitoring

Monitoring is an important part of any production system. There are many tools available and `pgBackRest` can be monitored on any of them with a little work.

pgBackRest can output information about the repository in JSON format which includes a list of all backups for each stanza and WAL archive info. A script is needed to extract information in a format that the monitoring system can understand.

pg-primary Get pgBackRest info in JSON format

```
sudo -u postgres pgbackrest --output=json info
```

```
[
  {
    "archive" : [
      {
        "database" : {
          "id" : 1
        },
        "id" : "10-1",
        "max" : "00000002000000000000000006",
        "min" : "00000001000000000000000002"
      }
    ],
    "backup" : [
      {
        "archive" : {
          "start" : "00000001000000000000000002",
          "stop" : "00000001000000000000000002"
        },
        "backrest" : {
          "format" : 5,
          "version" : "2.11"
        },
        "database" : {
          "id" : 1
        },
        "info" : {
          "delta" : 23521617,
          "repository" : {
            "delta" : 2790309,
            "size" : 2790309
          },
          "size" : 23521617
        },
        "label" : "20190310-121148F",
        "prior" : null,
        "reference" : null,
        "timestamp" : {
          "start" : 1552219908,
          "stop" : 1552219915
        },
        "type" : "full"
      },
      {
        "archive" : {
          "start" : "00000001000000000000000003",
          "stop" : "00000001000000000000000003"
        },
        "backrest" : {
          "format" : 5,
          "version" : "2.11"
        },
        "database" : {
          "id" : 1
        },
        "info" : {
          "delta" : 8429,
          "repository" : {
            "delta" : 517,
            "size" : 2790309
          }
        }
      }
    ]
  }
]
```

```

    },
    "size" : 23521617
  },
  "label" : "20190310-121148F_20190310-121156D",
  "prior" : "20190310-121148F",
  "reference" : [
    "20190310-121148F"
  ],
  "timestamp" : {
    "start" : 1552219916,
    "stop" : 1552219920
  },
  "type" : "diff"
},
{
  "archive" : {
    "start" : "00000002000000000000000005",
    "stop" : "00000002000000000000000005"
  },
  "backrest" : {
    "format" : 5,
    "version" : "2.11"
  },
  "database" : {
    "id" : 1
  },
  "info" : {
    "delta" : 8421,
    "repository" : {
      "delta" : 469,
      "size" : 2790309
    },
    "size" : 23521617
  },
  "label" : "20190310-121148F_20190310-121211I",
  "prior" : "20190310-121148F_20190310-121156D",
  "reference" : [
    "20190310-121148F",
    "20190310-121148F_20190310-121156D"
  ],
  "timestamp" : {
    "start" : 1552219931,
    "stop" : 1552219936
  },
  "type" : "incr"
},
{
  "archive" : {
    "start" : "00000002000000000000000006",
    "stop" : "00000002000000000000000006"
  },
  "backrest" : {
    "format" : 5,
    "version" : "2.11"
  },
  "database" : {
    "id" : 1
  },
  "info" : {
    "delta" : 8429,
    "repository" : {
      "delta" : 517,
      "size" : 2790309
    },
    "size" : 23521617
  }
}

```

```

    },
    "label" : "20190310-121148F_20190310-121218I",
    "prior" : "20190310-121148F_20190310-121211I",
    "reference" : [
        "20190310-121148F"
    ],
    "timestamp" : {
        "start" : 1552219938,
        "stop" : 1552219941
    },
    "type" : "incr"
}
],
"cipher" : "aes-256-cbc",
"db" : [
    {
        "id" : 1,
        "system-id" : 6666733617317675250,
        "version" : "10"
    }
],
"name" : "demo",
"status" : {
    "code" : 0,
    "message" : "ok"
}
}
]

```

## In PostgreSQL

The PostgreSQL COPY command allows pgBackRest info to be loaded into a table. The following example wraps that logic in a function that can be used to perform real-time queries.

pg-primary Load pgBackRest info function for PostgreSQL

```

sudo -u postgres cat \
    /var/lib/postgresql/pgbackrest/doc/example/pgsql-pgbackrest-info.sql

-- An example of monitoring pgBackRest from within PostgreSQL
--
-- Use copy to export data from the pgBackRest info command into the jsonb
-- type so it can be queried directly by PostgreSQL.
--
-- Create monitor schema
create schema monitor;
--
-- Get pgBackRest info in JSON format
create function monitor.pgbackrest_info()
    returns jsonb AS $$
declare
    data jsonb;
begin
    -- Create a temp table to hold the JSON data
    create temp table temp_pgbackrest_data (data jsonb);

    -- Copy data into the table directory from the pgBackRest into command
    copy temp_pgbackrest_data (data)
        from program
            'pgbackrest --output=json info | tr '\n' ' ' (format text);

select temp_pgbackrest_data.data
    into data
    from temp_pgbackrest_data;

```

```
drop table temp_pgbackrest_data;

return data;
end $$ language plpgsql;
```

```
sudo -u postgres psql -f \
    /var/lib/postgresql/pgbackrest/doc/example/pgsql-pgbackrest-info.sql
```

Now the `monitor.pgbackrest_info()` function can be used to determine the last successful backup time and archived WAL for a stanza.

pg-primary Query last successful backup time and archived WAL

```
sudo -u postgres cat \
    /var/lib/postgresql/pgbackrest/doc/example/pgsql-pgbackrest-query.sql
```

```
-- Get last successful backup for each stanza
--
-- Requires the monitor.pgbackrest_info function.
with stanza as
(
    select data->'name' as name,
           data->'backup'->(
               jsonb_array_length(data->'backup') - 1) as last_backup,
           data->'archive'->(
               jsonb_array_length(data->'archive') - 1) as current_archive
    from jsonb_array_elements(monitor.pgbackrest_info()) as data
)
select name,
       to_timestamp(
           (last_backup->'timestamp'->>'stop')::numeric) as last_successful_backup,
       current_archive->>'max' as last_archived_wal
from stanza;
```

```
sudo -u postgres psql -f \
    /var/lib/postgresql/pgbackrest/doc/example/pgsql-pgbackrest-query.sql
```

```
 name | last_successful_backup | last_archived_wal
-----+-----+-----
"demo" | 2019-03-10 12:12:21+00 | 00000002000000000000000006
(1 row)
```

## Using jq

jq is a command-line utility that can easily extract data from JSON.

pg-primary Install jq utility

```
sudo apt-get install jq
```

Now jq can be used to query the last successful backup time for a stanza.

pg-primary Query last successful backup time

```
sudo -u postgres pgbackrest --output=json --stanza=demo info | \
    jq '.[0] | .backup[-1] | .timestamp.stop'
```

```
1552219941
```

Or the last archived WAL.

pg-primary Query last archived WAL

```
sudo -u postgres pgbackrest --output=json --stanza=demo info | \
    jq '.[0] | .archive[-1] | .max'
```

```
"00000002000000000000000006"
```

NOTE:

This syntax requires jq v1.5 .

## Retention

Generally it is best to retain as many backups as possible to provide a greater window for Point-in-Time Recovery , but practical concerns such as disk space must also be considered. Retention options remove older backups once they are no longer needed.

## Full Backup Retention

Set `repo1-retention-full` to the number of full backups required. New backups must be completed before expiration will occur — that means if `repo1-retention-full=2` then there will be three full backups stored before the oldest one is expired.

```
pg-primary : /etc/pgbackrest/pgbackrest.conf  Configure repo1-retention-full
```

```
[demo]
```

```
pg1-path=/var/lib/postgresql/10/demo
```

```
[global]
```

```
repo1-cipher-pass=zWaf6XtpjIVZC5444yXB+cgFDF17MxGlgkZSaoPvTGirhPygu4jOKOXf9LO4vjfO
```

```
repo1-cipher-type=aes-256-cbc
```

```
repo1-path=/var/lib/pgbackrest
```

```
repo1-retention-full=2
```

```
start-fast=y
```

```
[global:archive-push]
```

```
compress-level=3
```

Backup `repo1-retention-full=2` but currently there is only one full backup so the next full backup to run will not expire any full backups.

```
pg-primary  Perform a full backup
```

```
sudo -u postgres pgbackrest --stanza=demo --type=full \  
    --log-level-console=detail backup
```

```
[filtered 951 lines of output]
```

```
P00  INFO: backup command end: completed successfully
```

```
P00  INFO: expire command begin
```

```
P00  DETAIL: archive retention on backup 20190310-121148F, archiveId = 10-1, start =  
00000001000000000000000002
```

```
P00  DETAIL: no archive to remove, archiveId = 10-1
```

```
P00  INFO: expire command end: completed successfully
```

Archive *is* expired because WAL segments were generated before the oldest backup. These are not useful for recovery — only WAL segments generated after a backup can be used to recover that backup.

```
pg-primary  Perform a full backup
```

```
sudo -u postgres pgbackrest --stanza=demo --type=full \  
    --log-level-console=info backup
```

```
[filtered 950 lines of output]
```

```
P00  INFO: backup command end: completed successfully
```

```
P00  INFO: expire command begin
```

```
P00  INFO: expire full backup set: 20190310-121148F, 20190310-121148F_20190310-121156D,  
20190310-121148F_20190310-121211I, 20190310-121148F_20190310-121218I
```

```
P00  INFO: remove expired backup 20190310-121148F_20190310-121218I
```

```
P00  INFO: remove expired backup 20190310-121148F_20190310-121211I
```

```
[filtered 2 lines of output]
```

The 20190310-121148F full backup is expired and archive retention is based on the 20190310-121228F which is now the oldest full backup.

## Differential Backup Retention

Set `repo1-retention-diff` to the number of differential backups required. Differentials only rely on the prior full backup so it is possible to create a “rolling” set of differentials for the last day or more. This allows quick restores to recent points-in-time but reduces overall space consumption.

```
pg-primary : /etc/pgbackrest/pgbackrest.conf  Configure repo1-retention-diff
```

```
[demo]
pg1-path=/var/lib/postgresql/10/demo
```

```
[global]
repo1-cipher-pass=zWaf6XtpjIVZC5444yXB+cgFDF17MxGlgkZSaoPvTGirhPygu4jOKOXf9LO4vjfO
repo1-cipher-type=aes-256-cbc
repo1-path=/var/lib/pgbackrest
repo1-retention-diff=1
repo1-retention-full=2
start-fast=y
```

```
[global:archive-push]
compress-level=3
```

Backup `repo1-retention-diff=1` so two differentials will need to be performed before one is expired. An incremental backup is added to demonstrate incremental expiration. Incremental backups cannot be expired independently — they are always expired with their related full or differential backup.

```
pg-primary  Perform differential and incremental backups
```

```
sudo -u postgres pgbackrest --stanza=demo --type=diff backup
```

```
sudo -u postgres pgbackrest --stanza=demo --type=incr backup
```

Now performing a differential backup will expire the previous differential and incremental backups leaving only one differential backup.

```
pg-primary  Perform a differential backup
```

```
sudo -u postgres pgbackrest --stanza=demo --type=diff \
  --log-level-console=info backup
```

```
[filtered 10 lines of output]
P00  INFO: backup command end: completed successfully
P00  INFO: expire command begin

P00  INFO: expire diff backup set: 20190310-121237F_20190310-121245D,
     20190310-121237F_20190310-121250I

P00  INFO: remove expired backup 20190310-121237F_20190310-121250I
P00  INFO: remove expired backup 20190310-121237F_20190310-121245D
```

## Archive Retention

Although `pgBackRest` automatically removes archived WAL segments when expiring backups (the default expires WAL for full backups based on the `repo1-retention-full` option), it may be useful to expire archive more aggressively to save disk space. Note that full backups are treated as differential backups for the purpose of differential archive retention.

Expiring archive will never remove WAL segments that are required to make a backup consistent. However, since Point-in-Time-Recovery (PITR) only works on a continuous WAL stream, care should be taken when aggressively expiring archive outside of the normal backup expiration process.

```
pg-primary : /etc/pgbackrest/pgbackrest.conf  Configure repo1-retention-diff
```

```
[demo]
pg1-path=/var/lib/postgresql/10/demo
```

```
[global]
repo1-cipher-pass=zWaf6XtpjIVZC5444yXB+cgFDF17MxGlgkZSaoPvTGirhPygu4jOKOXf9LO4vjfO
repo1-cipher-type=aes-256-cbc
repo1-path=/var/lib/pgbackrest
repo1-retention-diff=2
```

```
repo1-retention-full=2
```

```
start-fast=y
```

```
[global:archive-push]
```

```
compress-level=3
```

```
pg-primary Perform differential backup
```

```
sudo -u postgres pgbackrest --stanza=demo --type=diff \  
--log-level-console=info backup
```

```
[filtered 7 lines of output]
```

```
P00 INFO: execute non-exclusive pg_stop_backup() and wait for all WAL segments to archive
```

```
P00 INFO: backup stop archive = 00000002000000000000000010, lsn = 0/100000F8
```

```
P00 INFO: new backup label = 20190310-121237F_20190310-121259D
```

```
P00 INFO: backup command end: completed successfully
```

```
P00 INFO: expire command begin
```

```
pg-primary Expire archive
```

```
sudo -u postgres pgbackrest --stanza=demo --log-level-console=detail \  
--repo1-retention-archive-type=diff --repo1-retention-archive=1 expire
```

```
P00 INFO: expire command begin 2.11: --log-level-console=detail --log-level-stderr=off  
--no-log-timestamp --repo1-cipher-pass= --repo1-cipher-type=aes-256-cbc  
--repo1-path=/var/lib/pgbackrest --repo1-retention-archive=1  
--repo1-retention-archive-type=diff --repo1-retention-diff=2 --repo1-retention-full=2  
--stanza=demo
```

```
P00 DETAIL: archive retention on backup 20190310-121228F, archiveId = 10-1, start =  
00000002000000000000000008, stop = 00000002000000000000000008
```

```
P00 DETAIL: archive retention on backup 20190310-121237F, archiveId = 10-1, start =  
0000000200000000000000000A, stop = 0000000200000000000000000A
```

```
P00 DETAIL: archive retention on backup 20190310-121237F_20190310-121253D, archiveId = 10-1, start  
= 0000000200000000000000000D, stop = 0000000200000000000000000D
```

```
P00 DETAIL: archive retention on backup 20190310-121237F_20190310-121259D, archiveId = 10-1, start  
= 00000002000000000000000010
```

```
P00 DETAIL: remove archive: archiveId = 10-1, start = 00000002000000000000000009, stop =  
00000002000000000000000009
```

```
P00 DETAIL: remove archive: archiveId = 10-1, start = 0000000200000000000000000B, stop =  
0000000200000000000000000C
```

```
P00 DETAIL: remove archive: archiveId = 10-1, start = 0000000200000000000000000E, stop =  
0000000200000000000000000F
```

```
P00 INFO: expire command end: completed successfully
```

The 20190310-121237F\_20190310-121253D differential backup has archived WAL segments that must be retained to make the older backups consistent even though they cannot be played any further forward with PITR. WAL segments generated after 20190310-121237F\_20190310-121253D but before 20190310-121237F\_20190310-121259D are removed. WAL segments generated after the new backup 20190310-121237F\_20190310-121259D remain and can be used for PITR.

Since full backups are considered differential backups for the purpose of differential archive retention, if a full backup is now performed with the same settings, only the archive for that full backup is retained for PITR.

## Restore

The Restore section introduces additional restore command features.



## Delta Option

Restore a Backup in Quick Start required the database cluster directory to be cleaned before the restore could be performed. The delta option allows pgBackRest to automatically determine which files in the database cluster directory can be preserved and which ones need to be restored from the backup — it also *removes* files not present in the backup manifest so it will dispose of divergent changes. This is accomplished by calculating a [SHA-1](#) cryptographic hash for each file in the database cluster directory. If the SHA-1 hash does not match the hash stored in the backup then that file will be restored. This operation is very efficient when combined with the process-max option. Since the PostgreSQL server is shut down during the restore, a larger number of processes can be used than might be desirable during a backup when the PostgreSQL server is running.

pg-primary Stop the demo cluster, perform delta restore

```
sudo pg_ctlcluster 10 demo stop
```

```
sudo -u postgres pgbackrest --stanza=demo --delta \  
--log-level-console=detail restore
```

```
[filtered 760 lines of output]
```

```
P01 DETAIL: restore file /var/lib/postgresql/10/demo/base/12977/PG_VERSION - exists and matches  
backup (3B, 99%) checksum 4143d3a341877154d6e95211464e1df1015b74bd
```

```
P01 DETAIL: restore file /var/lib/postgresql/10/demo/base/1/PG_VERSION - exists and matches backup  
(3B, 99%) checksum 4143d3a341877154d6e95211464e1df1015b74bd
```

```
P01 DETAIL: restore file /var/lib/postgresql/10/demo/PG_VERSION - exists and matches backup (3B,  
100%) checksum 4143d3a341877154d6e95211464e1df1015b74bd
```

```
P01 DETAIL: restore file /var/lib/postgresql/10/demo/global/6100_vm - exists and is zero size (0B,  
100%)
```

```
P01 DETAIL: restore file /var/lib/postgresql/10/demo/global/6100 - exists and is zero size (0B,  
100%)
```

```
[filtered 203 lines of output]
```

```
P01 DETAIL: restore file /var/lib/postgresql/10/demo/base/1/12815 - exists and is zero size (0B,  
100%)
```

```
P00 INFO: write /var/lib/postgresql/10/demo/recovery.conf
```

```
P00 INFO: restore global/pg_control (performed last to ensure aborted restores cannot be started)
```

```
P00 INFO: restore command end: completed successfully
```

pg-primary Restart PostgreSQL

```
sudo pg_ctlcluster 10 demo start
```

## Restore Selected Databases

There may be cases where it is desirable to selectively restore specific databases from a cluster backup. This could be done for performance reasons or to move selected databases to a machine that does not have enough space to restore the entire cluster backup.

To demonstrate this feature two databases are created: test1 and test2. A fresh backup is run so pgBackRest is aware of the new databases.

pg-primary Create two test databases and perform a backup

```
sudo -u postgres psql -c "create database test1;"
```

```
CREATE DATABASE
```

```
sudo -u postgres psql -c "create database test2;"
```

```
CREATE DATABASE
```

```
sudo -u postgres pgbackrest --stanza=demo --type=incr backup
```

Each test database will be seeded with tables and data to demonstrate that recovery works with selective restore.

pg-primary Create a test table in each database

```
sudo -u postgres psql -c "create table test1_table (id int); \  
insert into test1_table (id) values (1);" test1
```

```
INSERT 0 1
```

```
sudo -u postgres psql -c "create table test2_table (id int); \  
insert into test2_table (id) values (2);" test2
```

```
INSERT 0 1
```

One of the main reasons to use selective restore is to save space. The size of the test1 database is shown here so it can be compared with the disk utilization after a selective restore.

pg-primary Show space used by test1 database

```
sudo -u postgres du -sh /var/lib/postgresql/10/demo/base/24576
```

```
7.5M /var/lib/postgresql/10/demo/base/24576
```

Stop the cluster and restore only the test2 database. Built-in databases ( template0 , template1 , and postgres ) are always restored.

pg-primary Restore from last backup including only the test2 database

```
sudo pg_ctlcluster 10 demo stop
```

```
sudo -u postgres pgbackrest --stanza=demo --delta \  
--db-include=test2 restore
```

```
sudo pg_ctlcluster 10 demo start
```

Once recovery is complete the test2 database will contain all previously created tables and data.

pg-primary Demonstrate that the test2 database was recovered

```
sudo -u postgres psql -c "select * from test2_table;" test2
```

```
id  
----  
2  
(1 row)
```

The test1 database, despite successful recovery, is not accessible. This is because the entire database was restored as sparse, zeroed files. PostgreSQL can successfully apply WAL on the zeroed files but the database as a whole will not be valid because key files contain no data. This is purposeful to prevent the database from being accidentally used when it might contain partial data that was applied during WAL replay.

pg-primary Attempting to connect to the test1 database will produce an error

```
sudo -u postgres psql -c "select * from test1_table;" test1
```

```
psql: FATAL: relation mapping file "base/24576/pg_filenode.map" contains invalid data
```

Since the test1 database is restored with sparse, zeroed files it will only require as much space as the amount of WAL that is written during recovery. While the amount of WAL generated during a backup and applied during recovery can be significant it will generally be a small fraction of the total database size, especially for large databases where this feature is most likely to be useful.

It is clear that the test1 database uses far less disk space during the selective restore than it would have if the entire database had been restored.

pg-primary Show space used by test1 database after recovery

```
sudo -u postgres du -sh /var/lib/postgresql/10/demo/base/24576
```

```
176K /var/lib/postgresql/10/demo/base/24576
```

At this point the only action that can be taken on the invalid test1 database is drop database . pgBackRest does not automatically drop the database since this cannot be done until recovery is complete and the cluster is accessible.

pg-primary Drop the test1 database

```
sudo -u postgres psql -c "drop database test1;"
```

```
DROP DATABASE
```

Now that the invalid test1 database has been dropped only the test2 and built-in databases remain.

pg-primary List remaining databases

```
sudo -u postgres psql -c "select oid, datname from pg_database order by oid;"
```

```
 oid | datname
-----+-----
      1 | template1
 12977 | template0
 12978 | postgres
```

```
24577 | test2
```

```
(4 rows)
```

## Point-in-Time Recovery

Restore a Backup in Quick Start performed default recovery, which is to play all the way to the end of the WAL stream. In the case of a hardware failure this is usually the best choice but for data corruption scenarios (whether machine or human in origin) Point-in-Time Recovery (PITR) is often more appropriate.

Point-in-Time Recovery (PITR) allows the WAL to be played from the last backup to a specified time, transaction id, or recovery point. For common recovery scenarios time-based recovery is arguably the most useful. A typical recovery scenario is to restore a table that was accidentally dropped or data that was accidentally deleted. Recovering a dropped table is more dramatic so that's the example given here but deleted data would be recovered in exactly the same way.

pg-primary Backup the demo cluster and create a table with very important data

```
sudo -u postgres pgbackrest --stanza=demo --type=diff backup
```

```
sudo -u postgres psql -c "begin; \
  create table important_table (message text); \
  insert into important_table values ('Important Data'); \
  commit; \
  select * from important_table;"
```

```
 message
-----
```

```
Important Data
```

```
(1 row)
```

It is important to represent the time as reckoned by PostgreSQL and to include timezone offsets. This reduces the possibility of unintended timezone conversions and an unexpected recovery result.

pg-primary Get the time from PostgreSQL

```
sudo -u postgres psql -Atc "select current_timestamp"
```

```
2019-03-10 12:13:43.83879+00
```

Now that the time has been recorded the table is dropped. In practice finding the exact time that the table was dropped is a lot harder than in this example. It may not be possible to find the exact time, but some forensic work should be able to get you close.

pg-primary Drop the important table

```
sudo -u postgres psql -c "begin; \
  drop table important_table; \
  commit; \
  select * from important_table;"
```

```
ERROR:  relation "important_table" does not exist
```

```
LINE 1: ...le important_table;      commit;      select * from important_...
                                         ^
```

Now the restore can be performed with time-based recovery to bring back the missing table.

pg-primary Stop PostgreSQL , restore the demo cluster to 2019-03-10 12:13:43.83879+00 , and display recovery.conf

```
sudo pg_ctlcluster 10 demo stop
```

```
sudo -u postgres pgbackrest --stanza=demo --delta \  
--type=time "--target=2019-03-10 12:13:43.83879+00" \  
--target-action=promote restore
```

```
sudo -u postgres cat /var/lib/postgresql/10/demo/recovery.conf
```

```
restore_command = 'pgbackrest --stanza=demo archive-get %f "%p"'
```

```
recovery_target_time = '2019-03-10 12:13:43.83879+00'
```

```
recovery_target_action = 'promote'
```

The recovery.conf file has been automatically generated by pgBackRest so PostgreSQL can be started immediately. Once PostgreSQL has finished recovery the table will exist again and can be queried.

pg-primary Start PostgreSQL and check that the important table exists

```
sudo pg_ctlcluster 10 demo start
```

```
sudo -u postgres psql -c "select * from important_table"
```

```
message
```

```
Important Data
```

```
(1 row)
```

The PostgreSQL log also contains valuable information. It will indicate the time and transaction where the recovery stopped and also give the time of the last transaction to be applied.

pg-primary Examine the PostgreSQL log output

```
sudo -u postgres cat /var/log/postgresql/postgresql-10-demo.log
```

```
[filtered 2 lines of output]
```

```
LOG: listening on Unix socket "/var/run/postgresql/.s.PGSQL.5432"
```

```
LOG: database system was interrupted; last known up at 2019-03-10 12:13:37 UTC
```

```
LOG: starting point-in-time recovery to 2019-03-10 12:13:43.83879+00
```

```
LOG: restored log file "00000004.history" from archive
```

```
LOG: restored log file "00000004000000000000000015" from archive
```

```
[filtered 2 lines of output]
```

```
LOG: database system is ready to accept read only connections
```

```
LOG: restored log file "00000004000000000000000016" from archive
```

```
LOG: recovery stopping before commit of transaction 564, time 2019-03-10 12:13:44.188755+00
```

```
LOG: redo done at 0/160207B0
```

```
LOG: last completed transaction was at log time 2019-03-10 12:13:43.488976+00
```

```
LOG: selected new timeline ID: 5
```

```
LOG: archive recovery complete
```

```
[filtered 3 lines of output]
```

This example was rigged to give the correct result. If a backup after the required time is chosen then PostgreSQL will not be able to recover the lost table. PostgreSQL can only play forward, not backward. To demonstrate this the important table must be dropped (again).

pg-primary Drop the important table (again)

```
sudo -u postgres psql -c "begin; \  
drop table important_table; \  
commit; \  
select * from important_table;"
```

```
ERROR: relation "important_table" does not exist
```

```
LINE 1: ...le important_table;      commit;      select * from important_...  
                                         ^
```

Now take a new backup and attempt recovery from the new backup.

pg-primary Perform a backup then attempt recovery from that backup

```
sudo -u postgres pgbackrest --stanza=demo --type=incr backup
```

```
sudo pg_ctlcluster 10 demo stop
```

```
sudo -u postgres pgbackrest --stanza=demo --delta \  
--type=time "--target=2019-03-10 12:13:43.83879+00" --target-action=promote restore
```

```
sudo pg_ctlcluster 10 demo start
```

```
sudo -u postgres psql -c "select * from important_table"
```

```
ERROR: relation "important_table" does not exist
```

```
LINE 1: select * from important_table  
                                         ^
```

Looking at the log output it's not obvious that recovery failed to restore the table. The key is to look for the presence of the "recovery stopping before..." and "last completed transaction..." log messages. If they are not present then the recovery to the specified point-in-time was not successful.

pg-primary Examine the PostgreSQL log output to discover the recovery was not successful

```
sudo -u postgres cat /var/log/postgresql/postgresql-10-demo.log
```

```
[filtered 2 lines of output]  
LOG: listening on Unix socket "/var/run/postgresql/.s.PGSQL.5432"  
LOG: database system was interrupted; last known up at 2019-03-10 12:13:57 UTC  
LOG: starting point-in-time recovery to 2019-03-10 12:13:43.83879+00  
LOG: restored log file "00000005.history" from archive  
LOG: restored log file "000000050000000000000000000017" from archive  
LOG: redo starts at 0/17000028  
LOG: consistent recovery state reached at 0/170000F8  
LOG: database system is ready to accept read only connections  
LOG: redo done at 0/170000F8  
[filtered 8 lines of output]
```

Using an earlier backup will allow PostgreSQL to play forward to the correct time. The info command can be used to find the next to last backup.

pg-primary Get backup info for the demo cluster

```
sudo -u postgres pgbackrest info
```

```
stanza: demo  
status: ok  
cipher: aes-256-cbc  
  
db (current)  
wal archive min/max (10-1): 00000002000000000000000008/000000050000000000000017  
  
full backup: 20190310-121228F
```

```
timestamp start/stop: 2019-03-10 12:12:28 / 2019-03-10 12:12:36
wal start/stop: 00000002000000000000000008 / 00000002000000000000000008
database size: 22.4MB, backup size: 22.4MB
repository size: 2.7MB, repository backup size: 2.7MB
```

```
full backup: 20190310-121237F
```

```
timestamp start/stop: 2019-03-10 12:12:37 / 2019-03-10 12:12:44
wal start/stop: 0000000200000000000000000A / 0000000200000000000000000A
database size: 22.4MB, backup size: 22.4MB
repository size: 2.7MB, repository backup size: 2.7MB
```

```
diff backup: 20190310-121237F_20190310-121259D
```

```
timestamp start/stop: 2019-03-10 12:12:59 / 2019-03-10 12:13:03
wal start/stop: 00000002000000000000000010 / 00000002000000000000000010
database size: 22.4MB, backup size: 8.2KB
repository size: 2.7MB, repository backup size: 519B
backup reference list: 20190310-121237F
```

```
incr backup: 20190310-121237F_20190310-121314I
```

```
timestamp start/stop: 2019-03-10 12:13:14 / 2019-03-10 12:13:22
wal start/stop: 00000003000000000000000012 / 00000003000000000000000012
database size: 37MB, backup size: 15MB
repository size: 4.4MB, repository backup size: 1.8MB
backup reference list: 20190310-121237F, 20190310-121237F_20190310-121259D
```

```
diff backup: 20190310-121237F_20190310-121336D
```

```
timestamp start/stop: 2019-03-10 12:13:36 / 2019-03-10 12:13:42
wal start/stop: 00000004000000000000000015 / 00000004000000000000000015
database size: 29.7MB, backup size: 7.8MB
repository size: 3.5MB, repository backup size: 948.7KB
backup reference list: 20190310-121237F
```

```
incr backup: 20190310-121237F_20190310-121356I
```

```
timestamp start/stop: 2019-03-10 12:13:56 / 2019-03-10 12:14:02
wal start/stop: 00000005000000000000000017 / 00000005000000000000000017
database size: 29.7MB, backup size: 2MB
repository size: 3.5MB, repository backup size: 218.3KB
```

```
backup reference list: 20190310-121237F, 20190310-121237F_20190310-121336D
```

The default behavior for restore is to use the last backup but an earlier backup can be specified with the `-set` option.

`pg-primary` Stop PostgreSQL, restore from the selected backup, and start PostgreSQL

```
sudo pg_ctlcluster 10 demo stop
```

```
sudo -u postgres pgbackrest --stanza=demo --delta \
--type=time "--target=2019-03-10 12:13:43.83879+00" \
--set=20190310-121237F_20190310-121336D --target-action=promote restore
```

```
sudo pg_ctlcluster 10 demo start
```

```
sudo -u postgres psql -c "select * from important_table"
```

```
message
```

```
-----
Important Data
```

```
(1 row)
```

Now the the log output will contain the expected “recovery stopping before...” and “last completed transaction...” messages showing that the recovery was successful.

`pg-primary` Examine the PostgreSQL log output for log messages indicating success

```
sudo -u postgres cat /var/log/postgresql/postgresql-10-demo.log
```

```
[filtered 2 lines of output]
```

```
LOG: listening on Unix socket "/var/run/postgresql/.s.PGSQL.5432"
```

```
LOG: database system was interrupted; last known up at 2019-03-10 12:13:37 UTC
```

```
LOG: starting point-in-time recovery to 2019-03-10 12:13:43.83879+00
```

```
LOG: restored log file "00000004.history" from archive
```

```
LOG: restored log file "00000004000000000000000015" from archive
```

```
[filtered 2 lines of output]
```

```
LOG: database system is ready to accept read only connections
```

```
LOG: restored log file "00000004000000000000000016" from archive
```

```
LOG: recovery stopping before commit of transaction 564, time 2019-03-10 12:13:44.188755+00
```

```
LOG: redo done at 0/160207B0
```

```
LOG: last completed transaction was at log time 2019-03-10 12:13:43.488976+00
```

```
LOG: restored log file "00000005.history" from archive
```

```
LOG: restored log file "00000006.history" from archive
```

```
[filtered 5 lines of output]
```

## S3 Support

pg-primary Install packages required for S3 support

```
sudo apt-get install libio-socket-ssl-perl libxml-libxml-perl
```

pgBackRest supports storing repositories in Amazon S3. The bucket used to store the repository must be created in advance — pgBackRest will not do it automatically. The repository can be located in the bucket root ( / ) but it's usually best to place it in a subpath so S3 logs or other data can also be stored in the bucket without conflicts.

pg-primary : /etc/pgbackrest/pgbackrest.conf Configure S3

```
[demo]
```

```
pg1-path=/var/lib/postgresql/10/demo
```

```
[global]
```

```
process-max=4
```

```
repo1-cipher-type=none
```

```
repo1-path=/demo-repo
```

```
repo1-retention-diff=2
```

```
repo1-retention-full=2
```

```
repo1-s3-bucket=demo-bucket
```

```
repo1-s3-endpoint=s3.amazonaws.com
```

```
repo1-s3-key=accessKey1
```

```
repo1-s3-key-secret=verySecretKey1
```

```
repo1-s3-region=us-east-1
```

```
repo1-type=s3
```

```
start-fast=y
```

```
[global:archive-push]
```

```
compress-level=3
```

NOTE:

The region and endpoint will need to be configured to where the bucket is located. The values given here are for the us-east-1 region.

A role should be created to run pgBackRest and the bucket permissions should be set as restrictively as possible. This sample policy will restrict all reads and writes to the bucket and repository path.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {
```

```

    "Effect": "Allow",
    "Action": [
        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::demo-bucket"
    ],
    "Condition": {
        "StringEquals": {
            "s3:prefix": [
                "",
                "demo-repo"
            ],
            "s3:delimiter": [
                "/"
            ]
        }
    }
},
{
    "Effect": "Allow",
    "Action": [
        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::demo-bucket"
    ],
    "Condition": {
        "StringLike": {
            "s3:prefix": [
                "demo-repo/*"
            ]
        }
    }
},
{
    "Effect": "Allow",
    "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:DeleteObject"
    ],
    "Resource": [
        "arn:aws:s3:::demo-bucket/demo-repo/*"
    ]
}
]
}

```

Commands are run exactly as if the repository were stored on a local disk.

pg-primary Create the stanza

```
sudo -u postgres pgbackrest --stanza=demo --log-level-console=info stanza-create
```

```

P00 INFO: stanza-create command begin 2.11: --log-level-console=info --log-level-stderr=off
--no-log-timestamp --pg1-path=/var/lib/postgresql/10/demo --repo1-cipher-type=none
--repo1-path=/demo-repo --repo1-s3-bucket=demo-bucket --repo1-s3-endpoint=s3.amazonaws.com
--repo1-s3-key= --repo1-s3-key-secret= --repo1-s3-region=us-east-1 --no-repo1-s3-verify-ssl
--repo1-type=s3 --stanza=demo

```

```
P00 INFO: stanza-create command end: completed successfully
```

File creation time in S3 is relatively slow so commands benefit by increasing process-max to parallelize file creation.

pg-primary Backup the demo cluster



```
sudo -u postgres pgbackrest --stanza=demo \  
--log-level-console=info backup
```

```
P00 INFO: backup command begin 2.11: --log-level-console=info --log-level-stderr=off  
--no-log-timestamp --pg1-path=/var/lib/postgresql/10/demo --process-max=4  
--repo1-cipher-type=none --repo1-path=/demo-repo --repo1-retention-diff=2  
--repo1-retention-full=2 --repo1-s3-bucket=demo-bucket --repo1-s3-endpoint=s3.amazonaws.com  
--repo1-s3-key= --repo1-s3-key-secret= --repo1-s3-region=us-east-1 --no-repo1-s3-verify-ssl  
--repo1-type=s3 --stanza=demo --start-fast
```

```
P00 WARN: no prior backup exists, incr backup has been changed to full
```

```
P00 INFO: execute non-exclusive pg_start_backup() with label "pgBackRest backup started at  
2019-03-10 12:14:34": backup begins after the requested immediate checkpoint completes
```

```
P00 INFO: backup start archive = 00000007000000000000000017, lsn = 0/17000028  
[filtered 1238 lines of output]
```

```
P02 INFO: backup file /var/lib/postgresql/10/demo/base/1/12820 (0B, 100%)
```

```
P04 INFO: backup file /var/lib/postgresql/10/demo/base/1/12815 (0B, 100%)
```

```
P00 INFO: full backup size = 29.7MB
```

```
P00 INFO: execute non-exclusive pg_stop_backup() and wait for all WAL segments to archive
```

```
P00 INFO: backup stop archive = 00000007000000000000000017, lsn = 0/17000130  
[filtered 4 lines of output]
```

## Delete a Stanza

The stanza-delete command removes data in the repository associated with a stanza.

WARNING:

Use this command with caution — it will permanently remove all backups and archives from the pgBackRest repository for the specified stanza.

To delete a stanza:

- Shut down the PostgreSQL cluster associated with the stanza (or use `-force` to override).
- Run the stop command on the repository host.
- Run the stanza-delete command on the repository host.

Once the command successfully completes, it is the responsibility of the user to remove the stanza from all pgBackRest configuration files.

```
pg-primary Stop PostgreSQL cluster to be removed
```

```
sudo pg_ctlcluster 10 demo stop
```

```
pg-primary Stop pgBackRest for the stanza
```

```
sudo -u postgres pgbackrest --stanza=demo --log-level-console=info stop
```

```
P00 INFO: stop command begin 2.11: --log-level-console=info --log-level-stderr=off  
--no-log-timestamp --repo1-cipher-type=none --repo1-path=/demo-repo  
--repo1-s3-bucket=demo-bucket --repo1-s3-endpoint=s3.amazonaws.com --repo1-s3-key=  
--repo1-s3-key-secret= --repo1-s3-region=us-east-1 --no-repo1-s3-verify-ssl --repo1-type=s3  
--stanza=demo
```

```
P00 INFO: stop command end: completed successfully
```

```
pg-primary Delete the stanza
```

```
sudo -u postgres pgbackrest --stanza=demo --log-level-console=info stanza-delete
```

```
P00 INFO: stanza-delete command begin 2.11: --log-level-console=info --log-level-stderr=off  
--no-log-timestamp --pg1-path=/var/lib/postgresql/10/demo --repo1-cipher-type=none  
--repo1-path=/demo-repo --repo1-s3-bucket=demo-bucket --repo1-s3-endpoint=s3.amazonaws.com  
--repo1-s3-key= --repo1-s3-key-secret= --repo1-s3-region=us-east-1 --no-repo1-s3-verify-ssl  
--repo1-type=s3 --stanza=demo
```

```
P00 INFO: stanza-delete command end: completed successfully
```

## Dedicated Repository Host

The configuration described in Quickstart is suitable for simple installations but for enterprise configurations it is more typical to have a dedicated repository host where the backups and WAL archive files are stored. This separates the backups and WAL archive from the database server so database host failures have less impact. It is still a good idea to employ traditional backup software to backup the repository host.

### Installation

A new host named repository is created to store the cluster backups.

The pgbackrest user is created to own the pgBackRest repository. Any user can own the repository but it is best not to use postgres (if it exists) to avoid confusion.

repository Create pgbackrest user

```
sudo adduser --disabled-password --gecos "" pgbackrest
```

pgBackRest needs to be installed from a package or installed manually as shown here.

repository Copy pgBackRest binary from build host

```
sudo scp build:/root/pgbackrest-release-2.11/src/pgbackrest /usr/bin
```

```
sudo chmod 755 /usr/bin/pgbackrest
```

pgBackRest contains embedded Perl which requires some additional modules.

repository Install required Perl packages

```
sudo apt-get install libdbd-pg-perl
```

Finally, pgBackRest requires log and configuration directories and a configuration file.

repository Create pgBackRest configuration file and directories

```
sudo mkdir -p -m 770 /var/log/pgbackrest
```

```
sudo chown pgbackrest:pgbackrest /var/log/pgbackrest
```

```
sudo mkdir -p /etc/pgbackrest
```

```
sudo mkdir -p /etc/pgbackrest/conf.d
```

```
sudo touch /etc/pgbackrest/pgbackrest.conf
```

```
sudo chmod 640 /etc/pgbackrest/pgbackrest.conf
```

```
sudo chown pgbackrest:pgbackrest /etc/pgbackrest/pgbackrest.conf
```

repository Create the pgBackRest repository

```
sudo mkdir -p /var/lib/pgbackrest
```

```
sudo chmod 750 /var/lib/pgbackrest
```

```
sudo chown pgbackrest:pgbackrest /var/lib/pgbackrest
```

### Setup Passwordless SSH

pgBackRest requires passwordless SSH to enable communication between the hosts.

repository Create repository host key pair

```
sudo -u pgbackrest mkdir -m 750 /home/pgbackrest/.ssh
```

```
sudo -u pgbackrest ssh-keygen -f /home/pgbackrest/.ssh/id_rsa \  
-t rsa -b 4096 -N ""
```

pg-primary Create pg-primary host key pair

```
sudo -u postgres mkdir -m 750 -p /var/lib/postgresql/.ssh
```

```
sudo -u postgres ssh-keygen -f /var/lib/postgresql/.ssh/id_rsa \  
-t rsa -b 4096 -N ""
```

Exchange keys between repository and pg-primary .

repository Copy pg-primary public key to repository

```
(echo -n 'no-agent-forwarding,no-X11-forwarding,no-port-forwarding,' && \  
echo -n 'command="/usr/bin/pgbackrest ${SSH_ORIGINAL_COMMAND##* }" ' && \  
sudo ssh root@pg-primary cat /var/lib/postgresql/.ssh/id_rsa.pub) | \  
sudo -u pgbackrest tee -a /home/pgbackrest/.ssh/authorized_keys
```

pg-primary Copy repository public key to pg-primary

```
(echo -n 'no-agent-forwarding,no-X11-forwarding,no-port-forwarding,' && \  
echo -n 'command="/usr/bin/pgbackrest ${SSH_ORIGINAL_COMMAND##* }" ' && \  
sudo ssh root@repository cat /home/pgbackrest/.ssh/id_rsa.pub) | \  
sudo -u postgres tee -a /var/lib/postgresql/.ssh/authorized_keys
```

Test that connections can be made from repository to pg-primary and vice versa.

repository Test connection from repository to pg-primary

```
sudo -u pgbackrest ssh postgres@pg-primary
```

pg-primary Test connection from pg-primary to repository

```
sudo -u postgres ssh pgbackrest@repository
```

NOTE:

ssh has been configured to only allow pgBackRest to be run via passwordless ssh. This enhances security in the event that one of the service accounts is hijacked.

## Configuration

The repository host must be configured with the pg-primary host/user and database path. The primary will be configured as db1 to allow a standby to be added later.

repository : /etc/pgbackrest/pgbackrest.conf Configure pg1-host / pg1-host-user and pg1-path

```
[demo]  
pg1-host=pg-primary  
pg1-path=/var/lib/postgresql/10/demo
```

```
[global]  
repo1-path=/var/lib/pgbackrest  
repo1-retention-full=2  
start-fast=y
```

The database host must be configured with the repository host/user. The default for the repo1-host-user option is pgbackrest . If the postgres user does restores on the repository host it is best not to also allow the postgres user to perform backups. However, the postgres user can read the repository directly if it is in the same group as the pgbackrest user.

pg-primary : /etc/pgbackrest/pgbackrest.conf Configure repo1-host / repo1-host-user

```
[demo]  
pg1-path=/var/lib/postgresql/10/demo
```

```
[global]  
log-level-file=detail  
repo1-host=repository
```

Commands are run the same as on a single host configuration except that some commands such as backup and expire are run from the repository host instead of the database host.

Create the stanza in the new repository.

repository Create the stanza

```
sudo -u pgbackrest pgbackrest --stanza=demo stanza-create
```

Check that the configuration is correct on both the database and repository hosts. More information about the check command can be found in [Check the Configuration](#) .

pg-primary Check the configuration

```
sudo -u postgres pgbackrest --stanza=demo check
```

repository Check the configuration

```
sudo -u pgbackrest pgbackrest --stanza=demo check
```

## Perform a Backup

To perform a backup of the PostgreSQL cluster run pgBackRest with the backup command on the repository host.

repository Backup the demo cluster

```
sudo -u pgbackrest pgbackrest --stanza=demo backup
```

```
P00 WARN: no prior backup exists, incr backup has been changed to full
```

Since a new repository was created on the repository host the warning about the incremental backup changing to a full backup was emitted.

## Restore a Backup

To perform a restore of the PostgreSQL cluster run pgBackRest with the restore command on the database host.

pg-primary Stop the demo cluster, restore, and restart PostgreSQL

```
sudo pg_ctlcluster 10 demo stop
```

```
sudo -u postgres pgbackrest --stanza=demo --delta restore
```

```
sudo pg_ctlcluster 10 demo start
```

A new backup must be performed due to the timeline switch.

repository Backup the demo cluster

```
sudo -u pgbackrest pgbackrest --stanza=demo backup
```

## Parallel Backup / Restore

pgBackRest offers parallel processing to improve performance of compression and transfer. The number of processes to be used for this feature is set using the `-process-max` option.

It is usually best not to use more than 25% of available CPUs for the backup command. Backups don't have to run that fast as long as they are performed regularly and the backup process should not impact database performance, if at all possible.

The restore command can and should use all available CPUs because during a restore the PostgreSQL cluster is shut down and there is generally no other important work being done on the host. If the host contains multiple clusters then that should be considered when setting restore parallelism.

repository Perform a backup with single process

```
sudo -u pgbackrest pgbackrest --stanza=demo --type=full backup
```

repository : /etc/pgbackrest/pgbackrest.conf Configure pgBackRest to use multiple backup processes

```
[demo]
pg1-host=pg-primary
pg1-path=/var/lib/postgresql/10/demo
```

```
[global]
process-max=3
repo1-path=/var/lib/pgbackrest
repo1-retention-full=2
start-fast=y
```

repository Perform a backup with multiple processes

```
sudo -u pgbackrest pgbackrest --stanza=demo --type=full backup
```

repository Get backup info for the demo cluster

```
sudo -u pgbackrest pgbackrest info
```

```
stanza: demo
  status: ok
  cipher: none

  db (current)
    wal archive min/max (10-1): 0000000800000000000000020/0000000800000000000000022

    full backup: 20190310-121613F
```

```
timestamp start/stop: 2019-03-10 12:16:13 / 2019-03-10 12:16:27
```

```
wal start/stop: 0000000800000000000000020 / 0000000800000000000000020
database size: 29.8MB, backup size: 29.8MB
repository size: 3.5MB, repository backup size: 3.5MB
```

```
full backup: 20190310-121629F
```

```
timestamp start/stop: 2019-03-10 12:16:29 / 2019-03-10 12:16:43
```

```
wal start/stop: 0000000800000000000000022 / 0000000800000000000000022
database size: 29.8MB, backup size: 29.8MB
repository size: 3.5MB, repository backup size: 3.5MB
```

The performance of the last backup should be improved by using multiple processes. For very small backups the difference may not be very apparent, but as the size of the database increases so will time savings.

## Starting and Stopping

Sometimes it is useful to prevent pgBackRest from running on a system. For example, when failing over from a primary to a standby it's best to prevent pgBackRest from running on the old primary in case PostgreSQL gets restarted or can't be completely killed. This will also prevent pgBackRest from running on cron .

pg-primary Stop the pgBackRest services

```
sudo -u postgres pgbackrest stop
```

New pgBackRest processes will no longer run.

repository Attempt a backup

```
sudo -u pgbackrest pgbackrest --stanza=demo backup
```

```
P00 ERROR: [062]: raised from remote process on 'pg-primary': stop file exists for all stanzas
```

Specify the `-force` option to terminate any pgBackRest process that are currently running. If pgBackRest is already stopped then stopping again will generate a warning.

pg-primary Stop the pgBackRest services again

```
sudo -u postgres pgbackrest stop
```

```
P00 WARN: stop file already exists for all stanzas
```

Start pgBackRest processes again with the start command.

pg-primary Start the pgBackRest services

```
sudo -u postgres pgbackrest start
```

It is also possible to stop pgBackRest for a single stanza.

pg-primary Stop pgBackRest services for the demo stanza

```
sudo -u postgres pgbackrest --stanza=demo stop
```

New pgBackRest processes for the specified stanza will no longer run.

repository Attempt a backup

```
sudo -u pgbackrest pgbackrest --stanza=demo backup
```

```
P00 ERROR: [062]: raised from remote process on 'pg-primary': stop file exists for stanza demo
```

The stanza must also be specified when starting the pgBackRest processes for a single stanza.

pg-primary Start the pgBackRest services for the demo stanza

```
sudo -u postgres pgbackrest --stanza=demo start
```

## Replication

Replication allows multiple copies of a PostgreSQL cluster (called standbys) to be created from a single primary. The standbys are useful for balancing reads and to provide redundancy in case the primary host fails.

## Installation

A new host named pg-standby is created to run the standby.

pgBackRest needs to be installed from a package or installed manually as shown here.

pg-standby Copy pgBackRest binary from build host

```
sudo scp build:/root/pgbackrest-release-2.11/src/pgbackrest /usr/bin
```

```
sudo chmod 755 /usr/bin/pgbackrest
```

pgBackRest contains embedded Perl which requires some additional modules.

pg-standby Install required Perl packages

```
sudo apt-get install libdbd-pg-perl
```

Finally, pgBackRest requires log and configuration directories and a configuration file.

pg-standby Create pgBackRest configuration file and directories

```
sudo mkdir -p -m 770 /var/log/pgbackrest
```

```
sudo chown postgres:postgres /var/log/pgbackrest
```

```
sudo mkdir -p /etc/pgbackrest
```

```
sudo mkdir -p /etc/pgbackrest/conf.d
```

```
sudo touch /etc/pgbackrest/pgbackrest.conf
```

```
sudo chmod 640 /etc/pgbackrest/pgbackrest.conf
```

```
sudo chown postgres:postgres /etc/pgbackrest/pgbackrest.conf
```

## Setup Passwordless SSH

pgBackRest requires passwordless SSH to enable communication between the hosts.

pg-standby Create pg-standby host key pair

```
sudo -u postgres mkdir -m 750 -p /var/lib/postgresql/.ssh
```

```
sudo -u postgres ssh-keygen -f /var/lib/postgresql/.ssh/id_rsa \  
-t rsa -b 4096 -N ""
```

Exchange keys between repository and pg-standby .

repository Copy pg-standby public key to repository

```
(echo -n 'no-agent-forwarding,no-X11-forwarding,no-port-forwarding,' && \  
echo -n 'command="/usr/bin/pgbackrest ${SSH_ORIGINAL_COMMAND#* }" ' && \  
sudo ssh root@pg-standby cat /var/lib/postgresql/.ssh/id_rsa.pub) | \  
sudo -u pgbackrest tee -a /home/pgbackrest/.ssh/authorized_keys
```

pg-standby Copy repository public key to pg-standby

```
(echo -n 'no-agent-forwarding,no-X11-forwarding,no-port-forwarding,' && \  
echo -n 'command="/usr/bin/pgbackrest ${SSH_ORIGINAL_COMMAND#* }" ' && \  
sudo ssh root@repository cat /home/pgbackrest/.ssh/id_rsa.pub) | \  
sudo -u postgres tee -a /var/lib/postgresql/.ssh/authorized_keys
```

Test that connections can be made from repository to pg-standby and vice versa.

repository Test connection from repository to pg-standby

```
sudo -u pgbackrest ssh postgres@pg-standby
```

pg-standby Test connection from pg-standby to repository

```
sudo -u postgres ssh pgbackrest@repository
```

## Hot Standby

A hot standby performs replication using the WAL archive and allows read-only queries.

pgBackRest configuration is very similar to pg-primary except that the `standby_mode` setting will be enabled to keep the cluster in recovery mode when the end of the WAL stream has been reached.

pg-standby : `/etc/pgbackrest/pgbackrest.conf` Configure pgBackRest on the standby

```
[demo]  
pg1-path=/var/lib/postgresql/10/demo  
recovery-option=standby_mode=on
```

```
[global]  
log-level-file=detail  
repo1-host=repository
```

The demo cluster must be created (even though it will be overwritten on restore) in order to create the PostgreSQL configuration files.

pg-standby Create demo cluster

```
sudo pg_createcluster 10 demo
```

Now the standby can be created with the restore command.

pg-standby Restore the demo standby cluster

```
sudo -u postgres pgbackrest --stanza=demo --delta restore
```

```
sudo -u postgres cat /var/lib/postgresql/10/demo/recovery.conf
```

```
standby_mode = 'on'  
restore_command = 'pgbackrest --stanza=demo archive-get %f "%p"'
```

### NOTE:

The `standby_mode` setting has been written into the `recovery.conf` file. Configuring recovery settings in pgBackRest means that the `recovery.conf` file does not need to be stored elsewhere since it will be properly recreated with each restore. The `-type=preserve` option can be used with the restore to leave the existing `recovery.conf` file in place if that behavior is preferred.

The `hot_standby` setting must be enabled before starting PostgreSQL to allow read-only connections on pg-standby . Otherwise, connection attempts will be refused. The rest of the configuration is in case the standby is promoted to a primary.

pg-standby : `/etc/postgresql/10/demo/postgresql.conf` Configure PostgreSQL

```
archive_command = 'pgbackrest --stanza=demo archive-push %p'  
archive_mode = on  
hot_standby = on  
log_filename = 'postgresql.log'  
log_line_prefix = ''
```

```
max_wal_senders = 3
wal_level = hot_standby
```

pg-standby Start PostgreSQL

```
sudo pg_ctlcluster 10 demo start
```

The PostgreSQL log gives valuable information about the recovery. Note especially that the cluster has entered standby mode and is ready to accept read-only connections.

pg-standby Examine the PostgreSQL log output for log messages indicating success

```
sudo -u postgres cat /var/log/postgresql/postgresql-10-demo.log
```

```
[filtered 3 lines of output]
LOG:  listening on Unix socket "/var/run/postgresql/.s.PGSQL.5432"
LOG:  database system was interrupted; last known up at 2019-03-10 12:16:30 UTC

LOG:  entering standby mode

LOG:  restored log file "00000008.history" from archive
LOG:  restored log file "000000080000000000000000022" from archive
LOG:  redo starts at 0/22000028
LOG:  consistent recovery state reached at 0/220000F8

LOG:  database system is ready to accept read only connections

LOG:  incomplete startup packet
```

An easy way to test that replication is properly configured is to create a table on pg-primary .

pg-primary Create a new table on the primary

```
sudo -u postgres psql -c " \
begin; \
create table replicated_table (message text); \
insert into replicated_table values ('Important Data'); \
commit; \
select * from replicated_table";
```

```
message
-----
```

```
Important Data
```

```
(1 row)
```

And then query the same table on pg-standby .

pg-standby Query new table on the standby

```
sudo -u postgres psql -c "select * from replicated_table;"
```

```
ERROR:  relation "replicated_table" does not exist
```

```
LINE 1: select * from replicated_table;
                        ^
```

So, what went wrong? Since PostgreSQL is pulling WAL segments from the archive to perform replication, changes won't be seen on the standby until the WAL segment that contains those changes is pushed from pg-primary .

This can be done manually by calling pg\_switch\_wal() which pushes the current WAL segment to the archive (a new WAL segment is created to contain further changes).

pg-primary Call pg\_switch\_wal()

```
sudo -u postgres psql -c "select *, current_timestamp from pg_switch_wal()";
```

```
pg_switch_wal |          current_timestamp
-----+-----
0/2302A958    | 2019-03-10 12:17:22.344397+00
(1 row)
```



Now after a short delay the table will appear on pg-standby .

pg-standby Now the new table exists on the standby (may require a few retries)

```
sudo -u postgres psql -c " \
    select *, current_timestamp from replicated_table"
```

```
message | current_timestamp
-----+-----
```

```
Important Data | 2019-03-10 12:17:27.56201+00
```

(1 row)

Check the standby configuration for access to the repository.

pg-standby Check the configuration

```
sudo -u postgres pgbackrest --stanza=demo --log-level-console=info check
```

```
P00 INFO: check command begin 2.11: --log-level-console=info --log-level-file=detail
--log-level-stderr=off --no-log-timestamp --pg1-path=/var/lib/postgresql/10/demo
--repo1-host=repository --stanza=demo
```

```
P00 INFO: switch wal cannot be performed on the standby, all other checks passed successfully
```

```
P00 INFO: check command end: completed successfully
```

## Streaming Replication

Instead of relying solely on the WAL archive, streaming replication makes a direct connection to the primary and applies changes as soon as they are made on the primary. This results in much less lag between the primary and standby.

Streaming replication requires a user with the replication privilege.

pg-primary Create replication user

```
sudo -u postgres psql -c " \
    create user replicator password 'jw8s0F4' replication";
```

```
CREATE ROLE
```

The pg\_hba.conf file must be updated to allow the standby to connect as the replication user. Be sure to replace the IP address below with the actual IP address of your pg-primary . A reload will be required after modifying the pg\_hba.conf file.

pg-primary Create pg\_hba.conf entry for replication user

```
sudo -u postgres sh -c 'echo \
    "host replication replicator 172.17.0.6/32 md5" \
    >> /etc/postgresql/10/demo/pg_hba.conf '
```

```
sudo pg_ctlcluster 10 demo reload
```

The standby needs to know how to contact the primary so the primary\_conninfo setting will be configured in pgBackRest .

pg-standby : /etc/pgbackrest/pgbackrest.conf Set primary\_conninfo

```
[demo]
pg1-path=/var/lib/postgresql/10/demo
recovery-option=standby_mode=on
recovery-option=primary_conninfo=host=172.17.0.4 port=5432 user=replicator
```

```
[global]
log-level-file=detail
repo1-host=repository
```

It is possible to configure a password in the primary\_conninfo setting but using a .pgpass file is more flexible and secure.

pg-standby Configure the replication password in the .pgpass file.

```
sudo -u postgres sh -c 'echo \  
"172.17.0.4:*:replication:replicator:jw8s0F4" \  
>> /var/lib/postgresql/.pgpass'
```

```
sudo -u postgres chmod 600 /var/lib/postgresql/.pgpass
```

Now the standby can be created with the restore command.

pg-standby Stop PostgreSQL and restore the demo standby cluster

```
sudo pg_ctlcluster 10 demo stop
```

```
sudo -u postgres pgbackrest --stanza=demo --delta restore
```

```
sudo -u postgres cat /var/lib/postgresql/10/demo/recovery.conf
```

```
primary_conninfo = 'host=172.17.0.4 port=5432 user=replicator '  
standby_mode = 'on'  
restore_command = 'pgbackrest --stanza=demo archive-get %f "%p"'
```

pg-standby Start PostgreSQL

```
sudo pg_ctlcluster 10 demo start
```

The PostgreSQL log will confirm that streaming replication has started.

pg-standby Examine the PostgreSQL log output for log messages indicating success

```
sudo -u postgres cat /var/log/postgresql/postgresql-10-demo.log
```

```
[filtered 11 lines of output]  
LOG:  restored log file "000000080000000000000000023" from archive  
LOG:  incomplete startup packet  
  
LOG:  started streaming WAL from primary at 0/24000000 on timeline 8
```

Now when a table is created on pg-primary it will appear on pg-standby quickly and without the need to call `pg_switch_wal()`.

pg-primary Create a new table on the primary

```
sudo -u postgres psql -c " \  
begin; \  
create table stream_table (message text); \  
insert into stream_table values ('Important Data'); \  
commit; \  
select *, current_timestamp from stream_table";
```

message	current_timestamp
Important Data	2019-03-10 12:17:42.903932+00

(1 row)

pg-standby Query table on the standby

```
sudo -u postgres psql -c " \  
select *, current_timestamp from stream_table"
```

message	current_timestamp
Important Data	2019-03-10 12:17:43.488868+00

(1 row)

## Asynchronous Archiving

Asynchronous archiving is enabled with the `archive-async` option. This option enables asynchronous operation for both the `archive-push` and `archive-get` commands.

A spool path is required. The commands will store transient data here but each command works quite a bit differently so spool path usage is described in detail in each section.

`pg-primary` Create the spool directory

```
sudo mkdir -p -m 750 /var/spool/pgbackrest
```

```
sudo chown postgres:postgres /var/spool/pgbackrest
```

`pg-standby` Create the spool directory

```
sudo mkdir -p -m 750 /var/spool/pgbackrest
```

```
sudo chown postgres:postgres /var/spool/pgbackrest
```

The spool path must be configured and asynchronous archiving enabled. Asynchronous archiving automatically confers some benefit by reducing the number of connections made to remote storage, but setting `process-max` can drastically improve performance by parallelizing operations. Be sure not to set `process-max` so high that it affects normal database operations.

`pg-primary` : `/etc/pgbackrest/pgbackrest.conf` Configure the spool path and asynchronous archiving

```
[demo]
pg1-path=/var/lib/postgresql/10/demo
```

```
[global]
archive-async=y
log-level-file=detail
repo1-host=repository
spool-path=/var/spool/pgbackrest
```

```
[global:archive-get]
process-max=2
```

```
[global:archive-push]
process-max=2
```

`pg-standby` : `/etc/pgbackrest/pgbackrest.conf` Configure the spool path and asynchronous archiving

```
[demo]
pg1-path=/var/lib/postgresql/10/demo
recovery-option=standby__mode=on
recovery-option=primary__conninfo=host=172.17.0.4 port=5432 user=replicator
```

```
[global]
archive-async=y
log-level-file=detail
repo1-host=repository
spool-path=/var/spool/pgbackrest
```

```
[global:archive-get]
process-max=2
```

```
[global:archive-push]
process-max=2
```

NOTE:

`process-max` is configured using command sections so that the option is not used by backup and restore. This also allows different values for `archive-push` and `archive-get` .

For demonstration purposes streaming replication will be broken to force PostgreSQL to get WAL using the `restore_command` .

`pg-primary` Break streaming replication by changing the replication password

```
sudo -u postgres psql -c "alter user replicator password 'bogus'"
```

```
ALTER ROLE
```

pg-standby Restart standby to break connection

```
sudo pg_ctlcluster 10 demo restart
```

## Archive Push

The asynchronous archive-push command offloads WAL archiving to a separate process (or processes) to improve throughput. It works by “looking ahead” to see which WAL segments are ready to be archived beyond the request that PostgreSQL is currently making via the `archive_command`. WAL segments are transferred to the archive directly from the `pg_xlog` / `pg_wal` directory and success is only returned by the `archive_command` when the WAL segment has been safely stored in the archive.

The pool path holds the current status of WAL archiving. Status files written into the pool directory are typically zero length and should consume a minimal amount of space (a few MB at most) and very little IO. All the information in this directory can be recreated so it is not necessary to preserve the pool directory if the cluster is moved to new hardware.

### IMPORTANT:

In the original implementation of asynchronous archiving, WAL segments were copied to the pool directory before compression and transfer. The new implementation copies WAL directly from the `pg_xlog` directory. If asynchronous archiving was utilized in v1.12 or prior, read the v1.13 release notes carefully before upgrading.

The `[stanza]-archive-push-async.log` file can be used to monitor the activity of the asynchronous process. A good way to test this is to quickly push a number of WAL segments.

pg-primary Test parallel asynchronous archiving

```
sudo -u postgres psql -c " \
select pg_create_restore_point('test async push'); select pg_switch_wal(); \
select pg_create_restore_point('test async push'); select pg_switch_wal(); \
select pg_create_restore_point('test async push'); select pg_switch_wal(); \
select pg_create_restore_point('test async push'); select pg_switch_wal(); \
select pg_create_restore_point('test async push'); select pg_switch_wal();"
```

```
sudo -u postgres pgbackrest --stanza=demo --log-level-console=info check
```

```
P00 INFO: check command begin 2.11: --log-level-console=info --log-level-file=detail
--log-level-stderr=off --no-log-timestamp --pg1-path=/var/lib/postgresql/10/demo
--repo1-host=repository --stanza=demo
```

```
P00 INFO: WAL segment 0000000800000000000000029 successfully stored in the archive at
'/var/lib/pgbackrest/archive/demo/10-1/0000000800000000/0000000800000000000000029-dbd7a9c8545e22f8d
```

```
P00 INFO: check command end: completed successfully
```

Now the log file will contain parallel, asynchronous activity.

pg-primary Check results in the log

```
sudo -u postgres cat /var/log/pgbackrest/demo-archive-push-async.log
```

```
-----PROCESS START-----
P00 INFO: archive-push command begin 2.11: [pg_wal/0000000800000000000000024] --archive-async
--log-level-console=off --log-level-file=detail --log-level-stderr=off --no-log-timestamp
--pg1-path=/var/lib/postgresql/10/demo --process-max=2 --repo1-host=repository
--spool-path=/var/spool/pgbackrest --stanza=demo
```

```
P00 INFO: push 2 WAL file(s) to archive: 0000000800000000000000024...000000080000000000000025
P01 DETAIL: pushed WAL file 0000000800000000000000024 to archive
P02 DETAIL: pushed WAL file 0000000800000000000000025 to archive
```

```
P00 INFO: archive-push command end: completed successfully
```

```
-----PROCESS START-----
P00 INFO: archive-push command begin 2.11: [pg_wal/0000000800000000000000026] --archive-async
--log-level-console=off --log-level-file=detail --log-level-stderr=off --no-log-timestamp
--pg1-path=/var/lib/postgresql/10/demo --process-max=2 --repo1-host=repository
--spool-path=/var/spool/pgbackrest --stanza=demo
```

```
P00 INFO: push 4 WAL file(s) to archive: 0000000800000000000000026...00000008000000000000029
P02 DETAIL: pushed WAL file 0000000800000000000000027 to archive
P01 DETAIL: pushed WAL file 0000000800000000000000026 to archive
P02 DETAIL: pushed WAL file 0000000800000000000000028 to archive
P01 DETAIL: pushed WAL file 0000000800000000000000029 to archive
```

```
P00 INFO: archive-push command end: completed successfully
```

## Archive Get

The asynchronous archive-get command maintains a local queue of WAL to improve throughput. If a WAL segment is not found in the queue it is fetched from the repository along with enough consecutive WAL to fill the queue. The maximum size of the queue is defined by archive-get-queue-max . Whenever the queue is less than half full more WAL will be fetched to fill it.

Asynchronous operation is most useful in environments that generate a lot of WAL or have a high latency connection to the repository storage (i.e., S3 ). In the case of a high latency connection it may be a good idea to increase process-max .

The [stanza]-archive-get-async.log file can be used to monitor the activity of the asynchronous process.

pg-standby Check results in the log

```
sudo -u postgres cat /var/log/pgbackrest/demo-archive-get-async.log
```

```
-----PROCESS START-----
P00 INFO: archive-get-async command begin 2.11: [0000000800000000000000022,
0000000800000000000000023, 0000000800000000000000024, 0000000800000000000000025,
0000000800000000000000026, 0000000800000000000000027, 0000000800000000000000028,
0000000800000000000000029] --log-level-console=off --log-level-file=detail
--log-level-stderr=off --no-log-timestamp --pg1-path=/var/lib/postgresql/10/demo
--process-max=2 --repo1-host=repository --spool-path=/var/spool/pgbackrest --stanza=demo
P00 INFO: get 8 WAL file(s) from archive: 0000000800000000000000022...00000008000000000000029
```

```
P00 DETAIL: found 0000000800000000000000023 in the archive
```

```
P00 DETAIL: unable to find 0000000800000000000000024 in the archive
```

```
P00 DETAIL: found 0000000800000000000000022 in the archive
```

```
P00 DETAIL: unable to find 0000000800000000000000025 in the archive
```

```
P00 DETAIL: unable to find 0000000800000000000000026 in the archive
```

```
[filtered 45 lines of output]
```

```
P00 INFO: archive-get-async command begin 2.11: [0000000800000000000000024,
0000000800000000000000025, 0000000800000000000000026, 0000000800000000000000027,
0000000800000000000000028, 0000000800000000000000029, 000000080000000000000002A,
000000080000000000000002B] --log-level-console=off --log-level-file=detail
--log-level-stderr=off --no-log-timestamp --pg1-path=/var/lib/postgresql/10/demo
--process-max=2 --repo1-host=repository --spool-path=/var/spool/pgbackrest --stanza=demo
P00 INFO: get 8 WAL file(s) from archive: 0000000800000000000000024...0000000800000000000002B
```

```
P00 DETAIL: found 0000000800000000000000025 in the archive
```

```
P00 DETAIL: found 0000000800000000000000024 in the archive
```

```
P00 DETAIL: found 0000000800000000000000027 in the archive
```

```
P00 DETAIL: found 0000000800000000000000026 in the archive
```

```
P00 DETAIL: unable to find 0000000800000000000000029 in the archive
```

```
P00 DETAIL: unable to find 000000080000000000000002A in the archive
```

```
P00 DETAIL: unable to find 000000080000000000000002B in the archive
```

```
P00 DETAIL: found 0000000800000000000000028 in the archive
```

```
-----PROCESS START-----
[filtered 5 lines of output]
P00 DETAIL: unable to find 000000080000000000000002C in the archive
P00 DETAIL: unable to find 000000080000000000000002D in the archive
```

```
P00 DETAIL: found 0000000800000000000000029 in the archive
```

```
-----PROCESS START-----  
[filtered 24 lines of output]
```

pg-primary Fix streaming replication by changing the replication password

```
sudo -u postgres psql -c "alter user replicator password 'jw8s0F4'"
```

```
ALTER ROLE
```

## Backup from a Standby

pgBackRest can perform backups on a standby instead of the primary. Standby backups require the pg-standby host to be configured and the backup-standby option enabled. If more than one standby is configured then the first running standby found will be used for the backup.

repository : /etc/pgbackrest/pgbackrest.conf Configure pg2-host / pg2-host-user and pg2-path

```
[demo]  
pg1-host=pg-primary  
pg1-path=/var/lib/postgresql/10/demo  
pg2-host=pg-standby  
pg2-path=/var/lib/postgresql/10/demo
```

```
[global]  
backup-standby=y  
process-max=3  
repo1-path=/var/lib/pgbackrest  
repo1-retention-full=2  
start-fast=y
```

Both the primary and standby databases are required to perform the backup, though the vast majority of the files will be copied from the standby to reduce load on the primary. The database hosts can be configured in any order. pgBackRest will automatically determine which is the primary and which is the standby.

repository Backup the demo cluster from pg2

```
sudo -u pgbackrest pgbackrest --stanza=demo --log-level-console=detail backup
```

```
[filtered 2 lines of output]
```

```
P00 INFO: execute non-exclusive pg_start_backup() with label "pgBackRest backup started at  
2019-03-10 12:18:07": backup begins after the requested immediate checkpoint completes  
P00 INFO: backup start archive = 00000008000000000000000002B, lsn = 0/2B000028
```

```
P00 INFO: wait for replay on the standby to reach 0/2B000028  
P00 INFO: replay on the standby reached 0/2B0000D0, checkpoint 0/2B000060  
P01 INFO: backup file pg-primary:/var/lib/postgresql/10/demo/global/pg_control (8KB, 0%)  
checksum 6faf2545f769bd65fb89293f8c3f8673bb60e78a  
P01 INFO: backup file pg-primary:/var/lib/postgresql/10/demo/pg_logical/replorigin_checkpoint  
(8B, 0%) checksum 347fc8f2df71bd4436e38bd1516ccd7ea0d46532
```

```
P02 INFO: backup file pg-standby:/var/lib/postgresql/10/demo/base/12978/2608 (440KB, 19%)  
checksum 729d02a898cc893fbee91feaa4d43831ac30a5af  
P02 INFO: backup file pg-standby:/var/lib/postgresql/10/demo/base/12978/2673 (320KB, 33%)  
checksum 7ba44db585de5f0a74b287a59ed4361ece8ab457  
[filtered 39 lines of output]
```

This incremental backup shows that most of the files are copied from the pg-standby host and only a few are copied from the pg-primary host.

pgBackRest creates a standby backup that is identical to a backup performed on the primary. It does this by starting/stopping the backup on the pg-primary host, copying only files that are replicated from the pg-standby host, then copying the remaining few files from the pg-primary host. This means that logs and statistics from the primary database will be included in the backup.

# Upgrading PostgreSQL

Immediately after upgrading PostgreSQL to a newer major version, the pg-path for all pgBackRest configurations must be set to the new database location and the stanza-upgrade run on the repository host. If the database is offline use the `-no-online` option.

The following instructions are not meant to be a comprehensive guide for upgrading PostgreSQL, rather they outline the general process for upgrading a primary and standby with the intent of demonstrating the steps required to reconfigure pgBackRest. It is recommended that a backup be taken prior to upgrading.

pg-primary Stop old cluster

```
sudo pg_ctlcluster 10 demo stop
```

Stop the old cluster on the standby since it will be restored from the newly upgraded cluster.

pg-standby Stop old cluster

```
sudo pg_ctlcluster 10 demo stop
```

Create the new cluster and perform upgrade.

pg-primary Create new cluster and perform the upgrade

```
sudo -u postgres /usr/lib/postgresql/11/bin/initdb \  
-D /var/lib/postgresql/11/demo -k -A peer
```

```
sudo pg_createcluster 11 demo
```

```
sudo -u postgres sh -c 'cd /var/lib/postgresql && \  
/usr/lib/postgresql/11/bin/pg_upgrade \  
--old-bindir=/usr/lib/postgresql/10/bin \  
--new-bindir=/usr/lib/postgresql/11/bin \  
--old-datadir=/var/lib/postgresql/10/demo \  
--new-datadir=/var/lib/postgresql/11/demo \  
--old-options=" -c config_file=/etc/postgresql/10/demo/postgresql.conf" \  
--new-options=" -c config_file=/etc/postgresql/11/demo/postgresql.conf"'
```

```
[filtered 68 lines of output]  
Creating script to delete old cluster ok
```

Upgrade Complete

```
-----  
Optimizer statistics are not transferred by pg_upgrade so,  
[filtered 4 lines of output]
```

Configure the new cluster settings and port.

pg-primary : `/etc/postgresql/11/demo/postgresql.conf` Configure PostgreSQL

```
archive_command = 'pgbackrest -stanza=demo archive-push %p'  
archive_mode = on  
listen_addresses = '*'  
log_line_prefix = ''  
max_wal_senders = 3  
port = 5432  
wal_level = hot_standby
```

Update the pgBackRest configuration on all systems to point to the new cluster.

pg-primary : `/etc/pgbackrest/pgbackrest.conf` Upgrade the pg1-path

```
[demo]  
pg1-path=/var/lib/postgresql/11/demo
```

```
[global]  
archive-async=y  
log-level-file=detail  
repol-host=repository  
spool-path=/var/spool/pgbackrest
```

```
[global:archive-get]
```

```
process-max=2
```

```
[global:archive-push]
```

```
process-max=2
```

```
pg-standby : /etc/pgbackrest/pgbackrest.conf Upgrade the pg-path
```

```
[demo]
```

```
pg1-path=/var/lib/postgresql/11/demo
```

```
recovery-option=standby_mode=on
```

```
recovery-option=primary_conninfo=host=172.17.0.4 port=5432 user=replicator
```

```
[global]
```

```
archive-async=y
```

```
log-level-file=detail
```

```
repo1-host=repository
```

```
spool-path=/var/spool/pgbackrest
```

```
[global:archive-get]
```

```
process-max=2
```

```
[global:archive-push]
```

```
process-max=2
```

```
repository : /etc/pgbackrest/pgbackrest.conf Upgrade pg1-path and pg2-path , disable backup from standby
```

```
[demo]
```

```
pg1-host=pg-primary
```

```
pg1-path=/var/lib/postgresql/11/demo
```

```
pg2-host=pg-standby
```

```
pg2-path=/var/lib/postgresql/11/demo
```

```
[global]
```

```
backup-standby=n
```

```
process-max=3
```

```
repo1-path=/var/lib/pgbackrest
```

```
repo1-retention-full=2
```

```
start-fast=y
```

```
pg-primary Copy hba configuration
```

```
sudo cp /etc/postgresql/10/demo/pg_hba.conf \  
/etc/postgresql/11/demo/pg_hba.conf
```

Before starting the new cluster, the stanza-upgrade command must be run on the server where the pgBackRest repository is located.

```
repository Upgrade the stanza
```

```
sudo -u pgbackrest pgbackrest --stanza=demo --no-online \  
--log-level-console=info stanza-upgrade
```

```
P00 INFO: stanza-upgrade command begin 2.11: --no-backup-standby --log-level-console=info  
--log-level-stderr=off --no-log-timestamp --no-online --pg1-host=pg-primary  
--pg2-host=pg-standby --pg1-path=/var/lib/postgresql/11/demo  
--pg2-path=/var/lib/postgresql/11/demo --repo1-path=/var/lib/pgbackrest --stanza=demo
```

```
P00 INFO: stanza-upgrade command end: completed successfully
```

Start the new cluster and confirm it is successfully installed.

```
pg-primary Start new cluster
```

```
sudo pg_ctlcluster 11 demo start
```

Test configuration using the check command.

```
pg-primary Check configuration
```

```
sudo -u postgres pg_lsclusters
```



Ver	Cluster	Port	Status	Owner	Data directory	Log file
10	demo	5432	down	postgres	/var/lib/postgresql/10/demo /var/log/postgresql/postgresql-10-demo.log	
11	demo	5432	online	postgres	/var/lib/postgresql/11/demo /var/log/postgresql/postgresql-11-demo.log	

```
sudo -u postgres pgbackrest --stanza=demo check
```

Remove the old cluster.

pg-primary Remove old cluster

```
sudo pg_dropcluster 10 demo
```

Install the new PostgreSQL binaries on the standby and create the cluster.

pg-standby Remove old cluster and create the new cluster

```
sudo pg_dropcluster 10 demo
```

```
sudo pg_createcluster 11 demo
```

Run the check on the repository host. The warning regarding the standby being down is expected since the standby cluster is down. Running this command demonstrates that the repository server is aware of the standby and is configured properly for the primary server.

repository Check configuration

```
sudo -u pgbackrest pgbackrest --stanza=demo check
```

Run a full backup on the new cluster and then restore the standby from the backup. The backup type will automatically be changed to full if incr or diff is requested.

repository Run a full backup

```
sudo -u pgbackrest pgbackrest --stanza=demo --type=full backup
```

pg-standby Restore the demo standby cluster

```
sudo -u postgres pgbackrest --stanza=demo --delta restore
```

pg-standby : /etc/postgresql/11/demo/postgresql.conf Configure PostgreSQL

hot\_standby = on

pg-standby Start PostgreSQL and check the pgBackRest configuration

```
sudo pg_ctlcluster 11 demo start
```

```
sudo -u postgres pgbackrest --stanza=demo check
```

Backup from standby can be enabled now that the standby is restored.

repository : /etc/pgbackrest/pgbackrest.conf Reenable backup from standby

[demo]

pg1-host=pg-primary

pg1-path=/var/lib/postgresql/11/demo

pg2-host=pg-standby

pg2-path=/var/lib/postgresql/11/demo

[global]

backup-standby=y

process-max=3

repo1-path=/var/lib/pgbackrest

repo1-retention-full=2

start-fast=y

— title: Command Reference draft: false —

## Introduction

Commands are used to execute the various pgBackRest functions. Here the command options are listed exhaustively, that is, each option applicable to a command is listed with that command even if it applies to one or more other commands. This includes all the options that may also be configured in `pgbackrest.conf`.

Non-boolean options configured in `pgbackrest.conf` can be reset to default on the command-line by using the `reset-` prefix. This feature may be used to restore a backup directly on a repository host. Normally, pgBackRest will error because it can see that the database host is remote and restores cannot be done remotely. By adding `-reset-pg1-host` on the command-line, pgBackRest will ignore the remote database host and restore locally. It may be necessary to pass a new `-pg-path` to force the restore to happen in a specific path, i.e. not the path used on the database host.

The `no-` prefix may be used to set a boolean option to false on the command-line.

Any option may be set in an environment variable using the `PGBACKREST_` prefix and the option name in all caps replacing `-` with `_`, e.g. `pg1-path` becomes `PGBACKREST_PG1_PATH` and `stanza` becomes `PGBACKREST_STANZA`. Boolean options are represented as they would be in a configuration file, e.g. `PGBACKREST_COMPRESS="n"`, and `reset-*` variants are not allowed. Options that can be specified multiple times on the command-line or in a config file can be represented by separating the values with colons, e.g. `PGBACKREST_DB_INCLUDE="db1:db2"`.

Command-line options override environment options which override config file options.

## Archive Get Command ( `archive-get` )

WAL segments are required for restoring a PostgreSQL cluster or maintaining a replica.

### Command Options

**Asynchronous Archiving Option ( `-archive-async` )** Push/get WAL segments asynchronously.

Enables asynchronous operation for the `archive-push` and `archive-get` commands.

Asynchronous operation is more efficient because it can reuse connections and take advantage of parallelism. See the `spool-path`, `archive-get-queue-max`, and `archive-push-queue-max` options for more information.

```
default: n
example: --archive-async
```

**Maximum Archive Get Queue Size Option ( `-archive-get-queue-max` )** Maximum size of the pgBackRest archive-get queue.

Specifies the maximum size of the archive-get queue when `archive-async` is enabled. The queue is stored in the `spool-path` and is used to speed providing WAL to PostgreSQL.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024.

```
default: 134217728
allowed: 0-4503599627370496
example: --archive-get-queue-max=1073741824
```

**Archive Timeout Option ( `-archive-timeout` )** Archive timeout.

Set maximum time, in seconds, to wait for each WAL segment to reach the pgBackRest archive repository. The timeout applies to the check and backup commands when waiting for WAL segments required for backup consistency to be archived.

```
default: 60
allowed: 0.1-86400
example: --archive-timeout=30
```

### General Options

**Buffer Size Option ( `-buffer-size` )** Buffer size for file operations.

Set the buffer size used for copy, compress, and uncompress functions. A maximum of 3 buffers will be in use at a time per process. An additional maximum of 256K per process may be used for zlib buffers.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024. For example, the case-insensitive value 32k (or 32KB) can be used instead of 32768.

Allowed values, in bytes, are 16384 , 32768 , 65536 , 131072 , 262144 , 524288 , 1048576 , 2097152 , 4194304 , 8388608 , and 16777216 .

```
default: 4194304
example: --buffer-size=32K
```

**SSH client command Option ( `-cmd-ssh` )** Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the ssh executable is not in \$PATH.

```
default: ssh
example: --cmd-ssh=/usr/bin/ssh
```

**Compress Option ( `-compress` )** Use gzip file compression.

Backup files are compatible with command-line gzip tools.

```
default: y
example: --no-compress
```

**Compress Level Option ( `-compress-level` )** Compression level for stored files.

Sets the zlib level to be used for file compression when compress=y .

```
default: 6
allowed: 0-9
example: --compress-level=9
```

**Network Compress Level Option ( `-compress-level-network` )** Compression level for network transfer when compress=n .

Sets the zlib level to be used for protocol compression when compress=n and the database cluster is not on the same host as the repository. Protocol compression is used to reduce network traffic but can be disabled by setting compress-level-network=0 . When compress=y the compress-level-network setting is ignored and compress-level is used instead so that the file is only compressed once. SSH compression is always disabled.

```
default: 3
allowed: 0-9
example: --compress-level-network=1
```

**Config Option ( `-config` )** pgBackRest configuration file.

Use this option to specify a different configuration file than the default.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --config=/conf/pgbackrest/pgbackrest.conf
```

**Config Include Path Option ( `-config-include-path` )** Path to additional pgBackRest configuration files.

Configuration files existing in the specified location with extension .conf will be concatenated with the pgBackRest configuration file, resulting in one configuration file.

```
default: /etc/pgbackrest/conf.d
example: --config-include-path=/conf/pgbackrest/conf.d
```

**Config Path Option ( `-config-path` )** Base path of pgBackRest configuration files.

This setting is used to override the default base path setting for the `-config` and `-config-include-path` options unless they are explicitly set on the command-line.

For example, passing only `-config-path=/conf/pgbackrest` results in the `-config` default being set to `/conf/pgbackrest/pgbackrest.conf` and the `-config-include-path` default being set to `/conf/pgbackrest/conf.d` .

```
default: /etc/pgbackrest
example: --config-path=/conf/pgbackrest
```

**Database Timeout Option ( `-db-timeout` )** Database query timeout.

Sets the timeout, in seconds, for queries against the database. This includes the `pg_start_backup()` and `pg_stop_backup()` functions which can each take a substantial amount of time. Because of this the timeout should be kept high unless you know that these functions will return quickly (i.e. if you have set `startfast=y` and you know that the database cluster will not generate many WAL segments during the backup).

```
default: 1800
allowed: 0.1-604800
example: --db-timeout=600
```

**Lock Path Option ( `-lock-path` )** Path where lock files are stored.

The lock path provides a location for `pgBackRest` to create lock files to prevent conflicting operations from being run concurrently.

```
default: /tmp/pgbackrest
example: --lock-path=/backup/db/lock
```

**Neutral Umask Option ( `-neutral-umask` )** Use a neutral umask.

Sets the umask to 0000 so modes in the repository are created in a sensible way. The default directory mode is 0750 and default file mode is 0640. The lock and log directories set the directory and file mode to 0770 and 0660 respectively.

To use the executing user's umask instead specify `neutral-umask=n` in the config file or `-no-neutral-umask` on the command line.

```
default: y
example: --no-neutral-umask
```

**Process Maximum Option ( `-process-max` )** Max processes to use for compress/transfer.

Each process will perform compression and transfer to make the command run faster, but don't set `process-max` so high that it impacts database performance.

```
default: 1
allowed: 1-96
example: --process-max=4
```

**Protocol Timeout Option ( `-protocol-timeout` )** Protocol timeout.

Sets the timeout, in seconds, that the local or remote process will wait for a new message to be received on the protocol layer. This prevents processes from waiting indefinitely for a message. The `protocol-timeout` option must be greater than the `db-timeout` option.

```
default: 1830
allowed: 0.1-604800
example: --protocol-timeout=630
```

**Spool Path Option ( `-spool-path` )** Path where transient data is stored.

This path is used to store data for the asynchronous `archive-push` and `archive-get` command.

The asynchronous `archive-push` command writes acknowledgements into the spool path when it has successfully stored WAL in the archive (and errors on failure) so the foreground process can quickly notify PostgreSQL. Acknowledgement files are very small (zero on success and a few hundred bytes on error).

The asynchronous `archive-push` process queues WAL in the spool path so it can be provided very quickly when PostgreSQL requests it. Moving files to PostgreSQL is most efficient when the spool path is on the same filesystem as `pg_xlog` / `pg_wal`.

The data stored in the spool path is not strictly temporary since it can and should survive a reboot. However, loss of the data in the spool path is not a problem. `pgBackRest` will simply recheck each WAL segment to ensure it is safely archived for `archive-push` and rebuild the queue for `archive-get`.

The spool path is intended to be located on a local Posix-compatible filesystem, not a remote filesystem such as NFS or CIFS.

```
default: /var/spool/pgbackrest
example: --spool-path=/backup/db/spool
```

**Stanza Option ( `-stanza` )** Defines the stanza.

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as `app` or `dw`, rather than the local cluster name, such as `main` or `prod`.

```
example: --stanza=main
```

## Log Options

**Console Log Level Option ( `-log-level-console` )** Level for console logging.

The following log levels are supported:

- `off` - No logging at all (not recommended)
- `error` - Log only errors
- `warn` - Log warnings and errors
- `info` - Log info, warnings, and errors
- `detail` - Log detail, info, warnings, and errors
- `debug` - Log debug, detail, info, warnings, and errors
- `trace` - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-console=error
```

**File Log Level Option ( `-log-level-file` )** Level for file logging.

The following log levels are supported:

- `off` - No logging at all (not recommended)
- `error` - Log only errors
- `warn` - Log warnings and errors
- `info` - Log info, warnings, and errors
- `detail` - Log detail, info, warnings, and errors
- `debug` - Log debug, detail, info, warnings, and errors
- `trace` - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
```

```
example: --log-level-file=debug
```

**Std Error Log Level Option ( `-log-level-stderr` )** Level for stderr logging.

Specifies which log levels will output to `stderr` rather than `stdout` (specified by `log-level-console`). The timestamp and process will not be output to `stderr`.

The following log levels are supported:

- `off` - No logging at all (not recommended)
- `error` - Log only errors
- `warn` - Log warnings and errors
- `info` - Log info, warnings, and errors
- `detail` - Log detail, info, warnings, and errors
- `debug` - Log debug, detail, info, warnings, and errors
- `trace` - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-stderr=error
```

**Log Path Option ( `-log-path` )** Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
example: --log-path=/backup/db/log
```

**Log Subprocesses Option ( `-log-subprocess` )** Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file` .

```
default: n
example: --log-subprocess
```

**Log Timestamp Option ( `-log-timestamp` )** Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
example: --no-log-timestamp
```

## Repository Options

**Repository Cipher Type Option ( `-repo-cipher-type` )** Cipher used to encrypt the repository.

The following repository types are supported:

- none - The repository is not encrypted
- aes-256-cbc - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. S3) supports encryption.

```
default: none
example: --repo1-cipher-type=aes-256-cbc
```

**Repository Host Option ( `-repo-host` )** Repository host when operating remotely via SSH.

Make sure that trusted SSH authentication is configured between the PostgreSQL host and the repository host.

When backing up and archiving to a locally mounted filesystem this setting is not required.

```
example: --repo1-host=repo1.domain.com
```

Deprecated Name: `backup-host`

**Repository Host Command Option ( `-repo-host-cmd` )** pgBackRest exe path on the repository host.

Required only if the path to pgbackrest is different on the local and repository hosts. If not defined, the repository host exe path will be set the same as the local exe path.

```
example: --repo1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: `backup-cmd`

**Repository Host Configuration Option ( `-repo-host-config` )** pgBackRest repository host configuration file.

Sets the location of the configuration file on the repository host. This is only required if the repository host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --repo1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: `backup-config`

**Repository Host Configuration Include Path Option ( `-repo-host-config-include-path` )** pgBackRest repository host configuration include path.

Sets the location of the configuration include path on the repository host. This is only required if the repository host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: --repo1-host-config-include-path=/conf/pgbackrest/conf.d
```

**Repository Host Configuration Path Option ( `-repo-host-config-path` )** pgBackRest repository host configuration path.

Sets the location of the configuration path on the repository host. This is only required if the repository host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --repo1-host-config-path=/conf/pgbackrest
```

**Repository Host Port Option ( `-repo-host-port` )** Repository host port when repo-host is set.

Use this option to specify a non-default port for the repository host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --repo1-host-port=25
```

Deprecated Name: backup-ssh-port

**Repository Host User Option ( `-repo-host-user` )** Repository host user when repo-host is set.

Defines the user that will be used for operations on the repository host. Preferably this is not the postgres user but rather some other user like pgbackrest . If PostgreSQL runs on the repository host the postgres user can be placed in the pgbackrest group so it has read permissions on the repository without being able to damage the contents accidentally.

```
default: pgbackrest
example: --repo1-host-user=repo-user
```

Deprecated Name: backup-user

**Repository Path Option ( `-repo-path` )** Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: --repo1-path=/backup/db/backrest
```

**S3 Repository Bucket Option ( `-repo-s3-bucket` )** S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting repo-path=/ but it is usually best to specify a prefix, such as /repo , so logs and other AWS generated content can also be stored in the bucket.

```
example: --repo1-s3-bucket=pg-backup
```

**S3 SSL CA File Option ( `-repo-s3-ca-file` )** S3 SSL CA File.

Use a CA file other than the system default.

```
example: --repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

### **S3 SSL CA Path Option ( `--repo-s3-ca-path` )** S3 SSL CA Path.

Use a CA path other than the system default.

```
example: --repo1-s3-ca-path=/etc/pki/tls/certs
```

### **S3 Repository Endpoint Option ( `--repo-s3-endpoint` )** S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: --repo1-s3-endpoint=s3.amazonaws.com
```

### **S3 Repository Host Option ( `--repo-s3-host` )** S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: --repo1-s3-host=127.0.0.1
```

### **S3 Repository Region Option ( `--repo-s3-region` )** S3 repository region.

The AWS region where the bucket was created.

```
example: --repo1-s3-region=us-east-1
```

### **S3 Repository Verify SSL Option ( `--repo-s3-verify-ssl` )** Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
```

```
example: --no-repo1-s3-verify-ssl
```

### **Repository Type Option ( `--repo-type` )** Type of storage used for the repository.

The following repository types are supported:

- cifs - Like posix , but disables links and directory fsyncs
- posix - Posix-compliant file systems
- s3 - AWS Simple Storage Service

```
default: posix
```

```
example: --repo1-type=cifs
```

## **Stanza Options**

### **PostgreSQL Path Option ( `--pg-path` )** PostgreSQL data directory.

This should be the same as the `data_directory` setting in `postgresql.conf` . Even though this value can be read from `postgresql.conf` or PostgreSQL it is prudent to set it in case those resources are not available during a restore or offline backup scenario.

The `pg-path` option is tested against the value reported by PostgreSQL on every online backup so it should always be current.

```
example: --pg1-path=/data/db
```

Deprecated Name: `db-path`

## **Archive Push Command ( `archive-push` )**

The WAL segment may be pushed immediately to the archive or stored locally depending on the value of `archive-async`



## Command Options

**Asynchronous Archiving Option ( `-archive-async` )** Push/get WAL segments asynchronously.

Enables asynchronous operation for the archive-push and archive-get commands.

Asynchronous operation is more efficient because it can reuse connections and take advantage of parallelism. See the `spool-path` , `archive-get-queue-max` , and `archive-push-queue-max` options for more information.

```
default: n
example: --archive-async
```

**Maximum Archive Push Queue Size Option ( `-archive-push-queue-max` )** Maximum size of the PostgreSQL archive queue.

After the limit is reached, the following will happen:

1. pgBackRest will notify PostgreSQL that the WAL was successfully archived, then **DROP IT** .
2. A warning will be output to the Postgres log.

If this occurs then the archive log stream will be interrupted and PITR will not be possible past that point. A new backup will be required to regain full restore capability.

In asynchronous mode the entire queue will be dropped to prevent spurts of WAL getting through before the queue limit is exceeded again.

The purpose of this feature is to prevent the log volume from filling up at which point Postgres will stop completely. Better to lose the backup than have PostgreSQL go down.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024.

```
allowed: 0-4503599627370496
example: --archive-push-queue-max=1GB
```

Deprecated Name: `archive-queue-max`

**Archive Timeout Option ( `-archive-timeout` )** Archive timeout.

Set maximum time, in seconds, to wait for each WAL segment to reach the pgBackRest archive repository. The timeout applies to the check and backup commands when waiting for WAL segments required for backup consistency to be archived.

```
default: 60
allowed: 0.1-86400
example: --archive-timeout=30
```

## General Options

**Buffer Size Option ( `-buffer-size` )** Buffer size for file operations.

Set the buffer size used for copy, compress, and uncompress functions. A maximum of 3 buffers will be in use at a time per process. An additional maximum of 256K per process may be used for zlib buffers.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024. For example, the case-insensitive value `32k` (or `32KB`) can be used instead of `32768`.

Allowed values, in bytes, are `16384` , `32768` , `65536` , `131072` , `262144` , `524288` , `1048576` , `2097152` , `4194304` , `8388608` , and `16777216` .

```
default: 4194304
example: --buffer-size=32K
```

**SSH client command Option ( `-cmd-ssh` )** Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the ssh executable is not in `$PATH`.

```
default: ssh
example: --cmd-ssh=/usr/bin/ssh
```

**Compress Option ( `-compress` )** Use gzip file compression.

Backup files are compatible with command-line gzip tools.

```
default: y
example: --no-compress
```

**Compress Level Option ( `-compress-level` )** Compression level for stored files.

Sets the zlib level to be used for file compression when `compress=y` .

```
default: 6
allowed: 0-9
example: --compress-level=9
```

**Network Compress Level Option ( `-compress-level-network` )** Compression level for network transfer when `compress=n` .

Sets the zlib level to be used for protocol compression when `compress=n` and the database cluster is not on the same host as the repository. Protocol compression is used to reduce network traffic but can be disabled by setting `compress-level-network=0` . When `compress=y` the `compress-level-network` setting is ignored and `compress-level` is used instead so that the file is only compressed once. SSH compression is always disabled.

```
default: 3
allowed: 0-9
example: --compress-level-network=1
```

**Config Option ( `-config` )** pgBackRest configuration file.

Use this option to specify a different configuration file than the default.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --config=/conf/pgbackrest/pgbackrest.conf
```

**Config Include Path Option ( `-config-include-path` )** Path to additional pgBackRest configuration files.

Configuration files existing in the specified location with extension `.conf` will be concatenated with the pgBackRest configuration file, resulting in one configuration file.

```
default: /etc/pgbackrest/conf.d
example: --config-include-path=/conf/pgbackrest/conf.d
```

**Config Path Option ( `-config-path` )** Base path of pgBackRest configuration files.

This setting is used to override the default base path setting for the `-config` and `-config-include-path` options unless they are explicitly set on the command-line.

For example, passing only `-config-path=/conf/pgbackrest` results in the `-config` default being set to `/conf/pgbackrest/pgbackrest.conf` and the `-config-include-path` default being set to `/conf/pgbackrest/conf.d` .

```
default: /etc/pgbackrest
example: --config-path=/conf/pgbackrest
```

**Database Timeout Option ( `-db-timeout` )** Database query timeout.

Sets the timeout, in seconds, for queries against the database. This includes the `pg_start_backup()` and `pg_stop_backup()` functions which can each take a substantial amount of time. Because of this the timeout should be kept high unless you know that these functions will return quickly (i.e. if you have set `startfast=y` and you know that the database cluster will not generate many WAL segments during the backup).

```
default: 1800
allowed: 0.1-604800
example: --db-timeout=600
```

**Lock Path Option ( `-lock-path` )** Path where lock files are stored.

The lock path provides a location for pgBackRest to create lock files to prevent conflicting operations from being run concurrently.

```
default: /tmp/pgbackrest
example: --lock-path=/backup/db/lock
```

**Neutral Umask Option ( `-neutral-umask` )** Use a neutral umask.

Sets the umask to 0000 so modes in the repository are created in a sensible way. The default directory mode is 0750 and default file mode is 0640. The lock and log directories set the directory and file mode to 0770 and 0660 respectively.

To use the executing user's umask instead specify `neutral-umask=n` in the config file or `-no-neutral-umask` on the command line.

```
default: y
example: --no-neutral-umask
```

**Process Maximum Option ( `-process-max` )** Max processes to use for compress/transfer.

Each process will perform compression and transfer to make the command run faster, but don't set `process-max` so high that it impacts database performance.

```
default: 1
allowed: 1-96
example: --process-max=4
```

**Protocol Timeout Option ( `-protocol-timeout` )** Protocol timeout.

Sets the timeout, in seconds, that the local or remote process will wait for a new message to be received on the protocol layer. This prevents processes from waiting indefinitely for a message. The `protocol-timeout` option must be greater than the `db-timeout` option.

```
default: 1830
allowed: 0.1-604800
example: --protocol-timeout=630
```

**Spool Path Option ( `-spool-path` )** Path where transient data is stored.

This path is used to store data for the asynchronous archive-push and archive-get command.

The asynchronous archive-push command writes acknowledgements into the spool path when it has successfully stored WAL in the archive (and errors on failure) so the foreground process can quickly notify PostgreSQL . Acknowledgement files are very small (zero on success and a few hundred bytes on error).

The asynchronous archive-push process queues WAL in the spool path so it can be provided very quickly when PostgreSQL requests it. Moving files to PostgreSQL is most efficient when the spool path is on the same filesystem as `pg_xlog` / `pg_wal` .

The data stored in the spool path is not strictly temporary since it can and should survive a reboot. However, loss of the data in the spool path is not a problem. pgBackRest will simply recheck each WAL segment to ensure it is safely archived for archive-push and rebuild the queue for archive-get .

The spool path is intended to be located on a local Posix-compatible filesystem, not a remote filesystem such as NFS or CIFS .

```
default: /var/spool/pgbackrest
example: --spool-path=/backup/db/spool
```

**Stanza Option ( `-stanza` )** Defines the stanza.

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as `app` or `dw`, rather than the local cluster name, such as `main` or `prod`.

```
example: --stanza=main
```

## Log Options

**Console Log Level Option ( `-log-level-console` )** Level for console logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-console=error
```

**File Log Level Option ( `-log-level-file` )** Level for file logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
```

```
example: --log-level-file=debug
```

**Std Error Log Level Option ( `-log-level-stderr` )** Level for stderr logging.

Specifies which log levels will output to stderr rather than stdout (specified by `log-level-console`). The timestamp and process will not be output to stderr.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-stderr=error
```

**Log Path Option ( `-log-path` )** Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
```

```
example: --log-path=/backup/db/log
```

**Log Subprocesses Option ( `-log-subprocess` )** Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file`.

```
default: n
example: --log-subprocess
```

**Log Timestamp Option ( `-log-timestamp` )** Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
example: --no-log-timestamp
```

## Repository Options

**Repository Cipher Type Option ( `-repo-cipher-type` )** Cipher used to encrypt the repository.

The following repository types are supported:

- none - The repository is not encrypted
- aes-256-cbc - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. S3) supports encryption.

```
default: none
example: --repo1-cipher-type=aes-256-cbc
```

**Repository Host Option ( `-repo-host` )** Repository host when operating remotely via SSH.

Make sure that trusted SSH authentication is configured between the PostgreSQL host and the repository host.

When backing up and archiving to a locally mounted filesystem this setting is not required.

```
example: --repo1-host=repo1.domain.com
```

Deprecated Name: `backup-host`

**Repository Host Command Option ( `-repo-host-cmd` )** `pgBackRest` exe path on the repository host.

Required only if the path to `pgbackrest` is different on the local and repository hosts. If not defined, the repository host exe path will be set the same as the local exe path.

```
example: --repo1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: `backup-cmd`

**Repository Host Configuration Option ( `-repo-host-config` )** `pgBackRest` repository host configuration file.

Sets the location of the configuration file on the repository host. This is only required if the repository host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --repo1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: `backup-config`

**Repository Host Configuration Include Path Option ( `-repo-host-config-include-path` )** `pgBackRest` repository host configuration include path.

Sets the location of the configuration include path on the repository host. This is only required if the repository host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: --repo1-host-config-include-path=/conf/pgbackrest/conf.d
```

**Repository Host Configuration Path Option ( `--repo-host-config-path` )** pgBackRest repository host configuration path.

Sets the location of the configuration path on the repository host. This is only required if the repository host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --repo1-host-config-path=/conf/pgbackrest
```

**Repository Host Port Option ( `--repo-host-port` )** Repository host port when `repo-host` is set.

Use this option to specify a non-default port for the repository host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --repo1-host-port=25
```

Deprecated Name: `backup-ssh-port`

**Repository Host User Option ( `--repo-host-user` )** Repository host user when `repo-host` is set.

Defines the user that will be used for operations on the repository host. Preferably this is not the postgres user but rather some other user like `pgbackrest`. If PostgreSQL runs on the repository host the postgres user can be placed in the `pgbackrest` group so it has read permissions on the repository without being able to damage the contents accidentally.

```
default: pgbackrest
example: --repo1-host-user=repo-user
```

Deprecated Name: `backup-user`

**Repository Path Option ( `--repo-path` )** Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: --repo1-path=/backup/db/backrest
```

**S3 Repository Bucket Option ( `--repo-s3-bucket` )** S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting `repo-path=` but it is usually best to specify a prefix, such as `/repo`, so logs and other AWS generated content can also be stored in the bucket.

```
example: --repo1-s3-bucket=pg-backup
```

**S3 SSL CA File Option ( `--repo-s3-ca-file` )** S3 SSL CA File.

Use a CA file other than the system default.

```
example: --repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

**S3 SSL CA Path Option ( `--repo-s3-ca-path` )** S3 SSL CA Path.

Use a CA path other than the system default.

```
example: --repo1-s3-ca-path=/etc/pki/tls/certs
```

**S3 Repository Endpoint Option ( `--repo-s3-endpoint` )** S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: --repo1-s3-endpoint=s3.amazonaws.com
```

**S3 Repository Host Option ( `-repo-s3-host` )** S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: --repo1-s3-host=127.0.0.1
```

**S3 Repository Region Option ( `-repo-s3-region` )** S3 repository region.

The AWS region where the bucket was created.

```
example: --repo1-s3-region=us-east-1
```

**S3 Repository Verify SSL Option ( `-repo-s3-verify-ssl` )** Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
```

```
example: --no-repo1-s3-verify-ssl
```

**Repository Type Option ( `-repo-type` )** Type of storage used for the repository.

The following repository types are supported:

- cifs - Like posix , but disables links and directory fsyncs
- posix - Posix-compliant file systems
- s3 - AWS Simple Storage Service

```
default: posix
```

```
example: --repo1-type=cifs
```

## Stanza Options

**PostgreSQL Path Option ( `-pg-path` )** PostgreSQL data directory.

This should be the same as the `data_directory` setting in `postgresql.conf` . Even though this value can be read from `postgresql.conf` or PostgreSQL it is prudent to set it in case those resources are not available during a restore or offline backup scenario.

The `pg-path` option is tested against the value reported by PostgreSQL on every online backup so it should always be current.

```
example: --pg1-path=/data/db
```

Deprecated Name: `db-path`

## Backup Command ( `backup` )

`pgBackRest` does not have a built-in scheduler so it's best to run it from cron or some other scheduling mechanism.

### Command Options

**Check Archive Option ( `-archive-check` )** Check that WAL segments are in the archive before backup completes.

Checks that all WAL segments required to make the backup consistent are present in the WAL archive. It's a good idea to leave this as the default unless you are using another method for archiving.

This option must be enabled if `archive-copy` is enabled.

```
default: y
```

```
example: --no-archive-check
```

**Copy Archive Option ( `-archive-copy` )** Copy WAL segments needed for consistency to the backup.

This slightly paranoid option protects against corruption in the WAL segment archive by storing the WAL segments required for consistency directly in the backup. WAL segments are still stored in the archive so this option will use additional space.

On restore, the WAL segments will be present in `pg_xlog/pg_wal` and PostgreSQL will use them in preference to calling the `restore_command`.

The `archive-check` option must be enabled if `archive-copy` is enabled.

```
default: n
example: --archive-copy
```

**Archive Timeout Option ( `-archive-timeout` )** Archive timeout.

Set maximum time, in seconds, to wait for each WAL segment to reach the pgBackRest archive repository. The timeout applies to the check and backup commands when waiting for WAL segments required for backup consistency to be archived.

```
default: 60
allowed: 0.1-86400
example: --archive-timeout=30
```

**Backup from Standby Option ( `-backup-standby` )** Backup from the standby cluster.

Enable backup from standby to reduce load on the primary cluster. This option requires that both the primary and standby hosts be configured.

```
default: n
example: --backup-standby
```

**Page Checksums Option ( `-checksum-page` )** Validate data page checksums.

Directs pgBackRest to validate all data page checksums while backing up a cluster. This option is automatically enabled when data page checksums are enabled on the cluster.

Failures in checksum validation will not abort a backup. Rather, warnings will be emitted in the log (and to the console with default settings) and the list of invalid pages will be stored in the backup manifest.

```
example: --no-checksum-page
```

**Path/File Exclusions Option ( `-exclude` )** Exclude paths/files from the backup.

All exclusions are relative to `$PGDATA`. If the exclusion ends with `/` then only files in the specified directory will be excluded, e.g. `-exclude=junk/` will exclude all files in the `$PGDATA/junk` directory but include the directory itself. If the exclusion does not end with `/` then the file may match the exclusion exactly or match with `/` appended to the exclusion, e.g. `-exclude=junk` will exclude the `$PGDATA/junk` directory and all the files it contains.

**Be careful using this feature – it is very easy to exclude something critical that will make the backup inconsistent. Be sure to test your restores!**

All excluded files will be logged at info level along with the exclusion rule. Be sure to audit the list of excluded files to ensure nothing unexpected is being excluded.

NOTE:

Exclusions are not honored on delta restores. Any files/directories that were excluded by the backup will be *removed* on delta restore.

This option should not be used to exclude PostgreSQL logs from a backup. Logs can be moved out of the `PGDATA` directory using the PostgreSQL `log_directory` setting, which has the benefit of allowing logs to be preserved after a restore.

Multiple exclusions may be specified on the command-line or in a configuration file.

```
example: --exclude=junk/
```



**Force Option ( `-force` )** Force an offline backup.

When used with `--no-start-stop` a backup will be run even if pgBackRest thinks that PostgreSQL is running. **This option should be used with extreme care as it will likely result in a bad backup.**

There are some scenarios where a backup might still be desirable under these conditions. For example, if a server crashes and the database cluster volume can only be mounted read-only, it would be a good idea to take a backup even if `postmaster.pid` is present. In this case it would be better to revert to the prior backup and replay WAL, but possibly there is a very important transaction in a WAL segment that did not get archived.

```
default: n
example: --force
```

**Manifest Save Threshold Option ( `--manifest-save-threshold` )** Manifest save threshold during backup.

Defines how often the manifest will be saved during a backup. Saving the manifest is important because it stores the checksums and allows the resume function to work efficiently. The actual threshold used is 1% of the backup size or `manifest-save-threshold`, whichever is greater.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024.

```
default: 1073741824
allowed: 1-1099511627776
example: --manifest-save-threshold=5G
```

**Online Option ( `--online` )** Perform an online backup.

Specifying `--no-online` prevents pgBackRest from running `pg_start_backup()` and `pg_stop_backup()` on the database cluster. In order for this to work PostgreSQL should be shut down and pgBackRest will generate an error if it is not.

The purpose of this option is to allow offline backups. The `pg_xlog / pg_wal` directory is copied as-is and `archive-check` is automatically disabled for the backup.

```
default: y
example: --no-online
```

**Resume Option ( `--resume` )** Allow resume of failed backup.

Defines whether the resume feature is enabled. Resume can greatly reduce the amount of time required to run a backup after a previous backup of the same type has failed. It adds complexity, however, so it may be desirable to disable in environments that do not require the feature.

```
default: y
example: --no-resume
```

**Start Fast Option ( `--start-fast` )** Force a checkpoint to start backup quickly.

Forces a checkpoint (by passing `y` to the `fast` parameter of `pg_start_backup()`) so the backup begins immediately. Otherwise the backup will start after the next regular checkpoint.

This feature only works in PostgreSQL  $\geq 8.4$ .

```
default: n
example: --start-fast
```

**Stop Auto Option ( `--stop-auto` )** Stop prior failed backup on new backup.

This will only be done if an exclusive advisory lock can be acquired to demonstrate that the prior failed backup process has really stopped.

This feature relies on `pg_is_in_backup()` so only works on PostgreSQL  $\geq 9.3$ .

The setting is disabled by default because it assumes that pgBackRest is the only process doing exclusive online backups. It depends on an advisory lock that only pgBackRest sets so it may abort other processes that do exclusive online backups. Note that `base_backup` and `pg_dump` are safe to use with this setting because they do not call `pg_start_backup()` so are not exclusive.

```
default: n
example: --stop-auto
```

**Type Option ( `-type` )** Backup type.

The following backup types are supported:

- full - all database cluster files will be copied and there will be no dependencies on previous backups.
- incr - incremental from the last successful backup.
- diff - like an incremental backup but always based on the last full backup.

```
default: incr
example: --type=full
```

## General Options

**Buffer Size Option ( `-buffer-size` )** Buffer size for file operations.

Set the buffer size used for copy, compress, and uncompress functions. A maximum of 3 buffers will be in use at a time per process. An additional maximum of 256K per process may be used for zlib buffers.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024. For example, the case-insensitive value 32k (or 32KB) can be used instead of 32768.

Allowed values, in bytes, are 16384 , 32768 , 65536 , 131072 , 262144 , 524288 , 1048576 , 2097152 , 4194304 , 8388608 , and 16777216 .

```
default: 4194304
example: --buffer-size=32K
```

**SSH client command Option ( `-cmd-ssh` )** Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the ssh executable is not in \$PATH.

```
default: ssh
example: --cmd-ssh=/usr/bin/ssh
```

**Compress Option ( `-compress` )** Use gzip file compression.

Backup files are compatible with command-line gzip tools.

```
default: y
example: --no-compress
```

**Compress Level Option ( `-compress-level` )** Compression level for stored files.

Sets the zlib level to be used for file compression when compress=y .

```
default: 6
allowed: 0-9
example: --compress-level=9
```

**Network Compress Level Option ( `-compress-level-network` )** Compression level for network transfer when compress=n .

Sets the zlib level to be used for protocol compression when compress=n and the database cluster is not on the same host as the repository. Protocol compression is used to reduce network traffic but can be disabled by setting compress-level-network=0 . When compress=y the compress-level-network setting is ignored and compress-level is used instead so that the file is only compressed once. SSH compression is always disabled.

```
default: 3
allowed: 0-9
example: --compress-level-network=1
```

**Config Option ( `-config` )** pgBackRest configuration file.

Use this option to specify a different configuration file than the default.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --config=/conf/pgbackrest/pgbackrest.conf
```

**Config Include Path Option ( `-config-include-path` )** Path to additional pgBackRest configuration files.

Configuration files existing in the specified location with extension `.conf` will be concatenated with the pgBackRest configuration file, resulting in one configuration file.

```
default: /etc/pgbackrest/conf.d
example: --config-include-path=/conf/pgbackrest/conf.d
```

**Config Path Option ( `-config-path` )** Base path of pgBackRest configuration files.

This setting is used to override the default base path setting for the `-config` and `-config-include-path` options unless they are explicitly set on the command-line.

For example, passing only `-config-path=/conf/pgbackrest` results in the `-config` default being set to `/conf/pgbackrest/pgbackrest.conf` and the `-config-include-path` default being set to `/conf/pgbackrest/conf.d`.

```
default: /etc/pgbackrest
example: --config-path=/conf/pgbackrest
```

**Database Timeout Option ( `-db-timeout` )** Database query timeout.

Sets the timeout, in seconds, for queries against the database. This includes the `pg_start_backup()` and `pg_stop_backup()` functions which can each take a substantial amount of time. Because of this the timeout should be kept high unless you know that these functions will return quickly (i.e. if you have set `startfast=y` and you know that the database cluster will not generate many WAL segments during the backup).

```
default: 1800
allowed: 0.1-604800
example: --db-timeout=600
```

**Delta Option ( `-delta` )** Restore or backup using checksums.

During a restore, by default the PostgreSQL data and tablespace directories are expected to be present but empty. This option performs a delta restore using checksums.

During a backup, this option will use checksums instead of the timestamps to determine if files will be copied.

```
default: n
example: --delta
```

**Lock Path Option ( `-lock-path` )** Path where lock files are stored.

The lock path provides a location for pgBackRest to create lock files to prevent conflicting operations from being run concurrently.

```
default: /tmp/pgbackrest
example: --lock-path=/backup/db/lock
```

**Neutral Umask Option ( `-neutral-umask` )** Use a neutral umask.

Sets the umask to 0000 so modes in the repository are created in a sensible way. The default directory mode is 0750 and default file mode is 0640. The lock and log directories set the directory and file mode to 0770 and 0660 respectively.

To use the executing user's umask instead specify `neutral-umask=n` in the config file or `-no-neutral-umask` on the command line.

```
default: y
example: --no-neutral-umask
```

**Process Maximum Option ( `-process-max` )** Max processes to use for compress/transfer.

Each process will perform compression and transfer to make the command run faster, but don't set process-max so high that it impacts database performance.

```
default: 1
allowed: 1-96
example: --process-max=4
```

**Protocol Timeout Option ( `-protocol-timeout` )** Protocol timeout.

Sets the timeout, in seconds, that the local or remote process will wait for a new message to be received on the protocol layer. This prevents processes from waiting indefinitely for a message. The protocol-timeout option must be greater than the db-timeout option.

```
default: 1830
allowed: 0.1-604800
example: --protocol-timeout=630
```

**Stanza Option ( `-stanza` )** Defines the stanza.

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as app or dw, rather than the local cluster name, such as main or prod.

```
example: --stanza=main
```

## Log Options

**Console Log Level Option ( `-log-level-console` )** Level for console logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
example: --log-level-console=error
```

**File Log Level Option ( `-log-level-file` )** Level for file logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
example: --log-level-file=debug
```

**Std Error Log Level Option ( `-log-level-stderr` )** Level for stderr logging.

Specifies which log levels will output to stderr rather than stdout (specified by `log-level-console`). The timestamp and process will not be output to stderr.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-stderr=error
```

**Log Path Option ( `-log-path` )** Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
```

```
example: --log-path=/backup/db/log
```

**Log Subprocesses Option ( `-log-subprocess` )** Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file`.

```
default: n
```

```
example: --log-subprocess
```

**Log Timestamp Option ( `-log-timestamp` )** Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
```

```
example: --no-log-timestamp
```

## Repository Options

**Repository Cipher Type Option ( `-repo-cipher-type` )** Cipher used to encrypt the repository.

The following repository types are supported:

- none - The repository is not encrypted
- aes-256-cbc - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. S3) supports encryption.

```
default: none
```

```
example: --repo1-cipher-type=aes-256-cbc
```

**Repository Hardlink Option ( `-repo-hardlink` )** Hardlink files between backups in the repository.

Enable hard-linking of files in differential and incremental backups to their full backups. This gives the appearance that each backup is a full backup at the file-system level. Be careful, though, because modifying files that are hard-linked can affect all the backups in the set.

```
default: n
```

```
example: --repo1-hardlink
```

Deprecated Name: `hardlink`

**Repository Path Option ( `--repo-path` )** Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: --repo1-path=/backup/db/backrest
```

**Archive Retention Option ( `--repo-retention-archive` )** Number of backups worth of continuous WAL to retain.

NOTE:

WAL segments required to make a backup consistent are always retained until the backup is expired regardless of how this option is configured.

If this value is not set, then the archive to expire will default to the `repo-retention-full` (or `repo-retention-diff`) value corresponding to the `repo-retention-archive-type` if set to `full` (or `diff`). This will ensure that WAL is only expired for backups that are already expired.

This option must be set if `repo-retention-archive-type` is set to `incr`. If disk space is at a premium, then this setting, in conjunction with `repo-retention-archive-type`, can be used to aggressively expire WAL segments. However, doing so negates the ability to perform PITR from the backups with expired WAL and is therefore **not** recommended.

```
allowed: 1-9999999
example: --repo1-retention-archive=2
```

Deprecated Name: `retention-archive`

**Archive Retention Type Option ( `--repo-retention-archive-type` )** Backup type for WAL retention.

If set to `full` pgBackRest will keep archive logs for the number of full backups defined by `repo-retention-archive`. If set to `diff` (differential) pgBackRest will keep archive logs for the number of full and differential backups defined by `repo-retention-archive`, meaning if the last backup taken was a full backup, it will be counted as a differential for the purpose of `repo-retention`. If set to `incr` (incremental) pgBackRest will keep archive logs for the number of full, differential, and incremental backups defined by `repo-retention-archive`. It is recommended that this setting not be changed from the default which will only expire WAL in conjunction with expiring full backups.

```
default: full
example: --repo1-retention-archive-type=diff
```

Deprecated Name: `retention-archive-type`

**Differential Retention Option ( `--repo-retention-diff` )** Number of differential backups to retain.

When a differential backup expires, all incremental backups associated with the differential backup will also expire. When not defined all differential backups will be kept until the full backups they depend on expire.

```
allowed: 1-9999999
example: --repo1-retention-diff=3
```

Deprecated Name: `retention-diff`

**Full Retention Option ( `--repo-retention-full` )** Number of full backups to retain.

When a full backup expires, all differential and incremental backups associated with the full backup will also expire. When the option is not defined a warning will be issued. If indefinite retention is desired then set the option to the max value.

```
allowed: 1-9999999
example: --repo1-retention-full=2
```

Deprecated Name: `retention-full`

**S3 Repository Bucket Option ( `--repo-s3-bucket` )** S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting `repo-path=` but it is usually best to specify a prefix, such as `/repo`, so logs and other AWS generated content can also be stored in the bucket.

```
example: --repo1-s3-bucket=pg-backup
```

### **S3 SSL CA File Option ( `--repo-s3-ca-file` )** S3 SSL CA File.

Use a CA file other than the system default.

```
example: --repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

### **S3 SSL CA Path Option ( `--repo-s3-ca-path` )** S3 SSL CA Path.

Use a CA path other than the system default.

```
example: --repo1-s3-ca-path=/etc/pki/tls/certs
```

### **S3 Repository Endpoint Option ( `--repo-s3-endpoint` )** S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: --repo1-s3-endpoint=s3.amazonaws.com
```

### **S3 Repository Host Option ( `--repo-s3-host` )** S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: --repo1-s3-host=127.0.0.1
```

### **S3 Repository Region Option ( `--repo-s3-region` )** S3 repository region.

The AWS region where the bucket was created.

```
example: --repo1-s3-region=us-east-1
```

### **S3 Repository Verify SSL Option ( `--repo-s3-verify-ssl` )** Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
```

```
example: --no-repo1-s3-verify-ssl
```

### **Repository Type Option ( `--repo-type` )** Type of storage used for the repository.

The following repository types are supported:

- cifs - Like posix , but disables links and directory fsyncs
- posix - Posix-compliant file systems
- s3 - AWS Simple Storage Service

```
default: posix
```

```
example: --repo1-type=cifs
```

## **Stanza Options**

### **PostgreSQL Host Option ( `--pg-host` )** PostgreSQL host for operating remotely via SSH.

Used for backups where the PostgreSQL host is different from the repository host.

```
example: --pg1-host=db.domain.com
```

Deprecated Name: db-host

**PostgreSQL Host Command Option ( `-pg-host-cmd` )** pgBackRest exe path on the PostgreSQL host.

Required only if the path to pgbackrest is different on the local and PostgreSQL hosts. If not defined, the database host exe path will be set the same as the local exe path.

```
example: --pg1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: db-cmd

**PostgreSQL Host Configuration Option ( `-pg-host-config` )** pgBackRest database host configuration file.

Sets the location of the configuration file on the PostgreSQL host. This is only required if the PostgreSQL host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
```

```
example: --pg1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: db-config

**PostgreSQL Host Configuration Include Path Option ( `-pg-host-config-include-path` )** pgBackRest database host configuration include path.

Sets the location of the configuration include path on the PostgreSQL host. This is only required if the PostgreSQL host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
```

```
example: --pg1-host-config-include-path=/conf/pgbackrest/conf.d
```

**PostgreSQL Host Configuration Path Option ( `-pg-host-config-path` )** pgBackRest database host configuration path.

Sets the location of the configuration path on the PostgreSQL host. This is only required if the PostgreSQL host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
```

```
example: --pg1-host-config-path=/conf/pgbackrest
```

**PostgreSQL Host Port Option ( `-pg-host-port` )** PostgreSQL host port when pg-host is set.

Use this option to specify a non-default port for the PostgreSQL host protocol. Currently only SSH is supported

```
allowed: 0-65535
```

```
example: --pg1-host-port=25
```

Deprecated Name: db-ssh-port

**PostgreSQL Host User Option ( `-pg-host-user` )** PostgreSQL host logon user when pg-host is set.

This user will also own the remote pgBackRest process and will initiate connections to PostgreSQL . For this to work correctly the user should be the PostgreSQL database cluster owner which is generally postgres , the default.

```
default: postgres
```

```
example: --pg1-host-user=db_owner
```

Deprecated Name: db-user

**PostgreSQL Path Option ( `-pg-path` )** PostgreSQL data directory.

This should be the same as the data\_directory setting in postgresql.conf . Even though this value can be read from postgresql.conf or PostgreSQL it is prudent to set it in case those resources are not available during a restore or offline backup scenario.

The pg-path option is tested against the value reported by PostgreSQL on every online backup so it should always be current.

```
example: --pg1-path=/data/db
```

Deprecated Name: db-path



**PostgreSQL Port Option ( `-pg-port` )** PostgreSQL port.

Port that PostgreSQL is running on. This usually does not need to be specified as most PostgreSQL clusters run on the default port.

```
default: 5432
allowed: 0-65535
example: --pg1-port=6543
```

Deprecated Name: db-port

**PostgreSQL Socket Path Option ( `-pg-socket-path` )** PostgreSQL unix socket path.

The unix socket directory that was specified when PostgreSQL was started. pgBackRest will automatically look in the standard location for your OS so there is usually no need to specify this setting unless the socket directory was explicitly modified with the `unix_socket_directory` setting in `postgres.conf`.

```
allowed: 0-65535
example: --pg1-socket-path=/var/run/postgresql
```

Deprecated Name: db-socket-path

## Check Command ( `check` )

The check command validates that pgBackRest and the `archive_command` setting are configured correctly for archiving and backups. It detects misconfigurations, particularly in archiving, that result in incomplete backups because required WAL segments did not reach the archive. The command can be run on the database or the repository host. The command may also be run on the standby host, however, since `pg_switch_xlog()` / `pg_switch_wal()` cannot be performed on the standby, the command will only test the repository configuration.

Note that `pg_create_restore_point('pgBackRest Archive Check')` and `pg_switch_xlog()` / `pg_switch_wal()` are called to force PostgreSQL to archive a WAL segment. Restore points are only supported in PostgreSQL  $\geq 9.1$  so for older versions the check command may fail if there has been no write activity since the last log rotation, therefore it is recommended that activity be generated by the user if there have been no writes since the last WAL switch before running the check command.

## Command Options

**Check Archive Option ( `-archive-check` )** Check that WAL segments are in the archive before backup completes.

Checks that all WAL segments required to make the backup consistent are present in the WAL archive. It's a good idea to leave this as the default unless you are using another method for archiving.

This option must be enabled if `archive-copy` is enabled.

```
default: y
example: --no-archive-check
```

**Archive Timeout Option ( `-archive-timeout` )** Archive timeout.

Set maximum time, in seconds, to wait for each WAL segment to reach the pgBackRest archive repository. The timeout applies to the check and backup commands when waiting for WAL segments required for backup consistency to be archived.

```
default: 60
allowed: 0.1-86400
example: --archive-timeout=30
```

**Backup from Standby Option ( `-backup-standby` )** Backup from the standby cluster.

Enable backup from standby to reduce load on the primary cluster. This option requires that both the primary and standby hosts be configured.

```
default: n
example: --backup-standby
```

## General Options

**Buffer Size Option ( `-buffer-size` )** Buffer size for file operations.

Set the buffer size used for copy, compress, and uncompress functions. A maximum of 3 buffers will be in use at a time per process. An additional maximum of 256K per process may be used for zlib buffers.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024. For example, the case-insensitive value 32k (or 32KB) can be used instead of 32768.

Allowed values, in bytes, are 16384 , 32768 , 65536 , 131072 , 262144 , 524288 , 1048576 , 2097152 , 4194304 , 8388608 , and 16777216 .

```
default: 4194304
```

```
example: --buffer-size=32K
```

**SSH client command Option ( `-cmd-ssh` )** Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the ssh executable is not in \$PATH.

```
default: ssh
```

```
example: --cmd-ssh=/usr/bin/ssh
```

**Compress Level Option ( `-compress-level` )** Compression level for stored files.

Sets the zlib level to be used for file compression when `compress=y` .

```
default: 6
```

```
allowed: 0-9
```

```
example: --compress-level=9
```

**Network Compress Level Option ( `-compress-level-network` )** Compression level for network transfer when `compress=n` .

Sets the zlib level to be used for protocol compression when `compress=n` and the database cluster is not on the same host as the repository. Protocol compression is used to reduce network traffic but can be disabled by setting `compress-level-network=0` . When `compress=y` the `compress-level-network` setting is ignored and `compress-level` is used instead so that the file is only compressed once. SSH compression is always disabled.

```
default: 3
```

```
allowed: 0-9
```

```
example: --compress-level-network=1
```

**Config Option ( `-config` )** pgBackRest configuration file.

Use this option to specify a different configuration file than the default.

```
default: /etc/pgbackrest/pgbackrest.conf
```

```
example: --config=/conf/pgbackrest/pgbackrest.conf
```

**Config Include Path Option ( `-config-include-path` )** Path to additional pgBackRest configuration files.

Configuration files existing in the specified location with extension `.conf` will be concatenated with the pgBackRest configuration file, resulting in one configuration file.

```
default: /etc/pgbackrest/conf.d
```

```
example: --config-include-path=/conf/pgbackrest/conf.d
```

**Config Path Option ( `-config-path` )** Base path of pgBackRest configuration files.

This setting is used to override the default base path setting for the `-config` and `-config-include-path` options unless they are explicitly set on the command-line.

For example, passing only `-config-path=/conf/pgbackrest` results in the `-config` default being set to `/conf/pgbackrest/pgbackrest.conf` and the `-config-include-path` default being set to `/conf/pgbackrest/conf.d` .

```
default: /etc/pgbackrest
```

```
example: --config-path=/conf/pgbackrest
```

**Database Timeout Option ( `-db-timeout` )** Database query timeout.

Sets the timeout, in seconds, for queries against the database. This includes the `pg_start_backup()` and `pg_stop_backup()` functions which can each take a substantial amount of time. Because of this the timeout should be kept high unless you know that these functions will return quickly (i.e. if you have set `startfast=y` and you know that the database cluster will not generate many WAL segments during the backup).

```
default: 1800
allowed: 0.1-604800
example: --db-timeout=600
```

**Neutral Umask Option ( `-neutral-umask` )** Use a neutral umask.

Sets the umask to 0000 so modes in the repository are created in a sensible way. The default directory mode is 0750 and default file mode is 0640. The lock and log directories set the directory and file mode to 0770 and 0660 respectively.

To use the executing user's umask instead specify `neutral-umask=n` in the config file or `-no-neutral-umask` on the command line.

```
default: y
example: --no-neutral-umask
```

**Protocol Timeout Option ( `-protocol-timeout` )** Protocol timeout.

Sets the timeout, in seconds, that the local or remote process will wait for a new message to be received on the protocol layer. This prevents processes from waiting indefinitely for a message. The `protocol-timeout` option must be greater than the `db-timeout` option.

```
default: 1830
allowed: 0.1-604800
example: --protocol-timeout=630
```

**Stanza Option ( `-stanza` )** Defines the stanza.

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as `app` or `dw`, rather than the local cluster name, such as `main` or `prod`.

```
example: --stanza=main
```

## Log Options

**Console Log Level Option ( `-log-level-console` )** Level for console logging.

The following log levels are supported:

- `off` - No logging at all (not recommended)
- `error` - Log only errors
- `warn` - Log warnings and errors
- `info` - Log info, warnings, and errors
- `detail` - Log detail, info, warnings, and errors
- `debug` - Log debug, detail, info, warnings, and errors
- `trace` - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
example: --log-level-console=error
```

**File Log Level Option ( `-log-level-file` )** Level for file logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
```

```
example: --log-level-file=debug
```

**Std Error Log Level Option ( `-log-level-stderr` )** Level for stderr logging.

Specifies which log levels will output to stderr rather than stdout (specified by `log-level-console` ). The timestamp and process will not be output to stderr .

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-stderr=error
```

**Log Path Option ( `-log-path` )** Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
```

```
example: --log-path=/backup/db/log
```

**Log Subprocesses Option ( `-log-subprocess` )** Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file` .

```
default: n
```

```
example: --log-subprocess
```

**Log Timestamp Option ( `-log-timestamp` )** Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
```

```
example: --no-log-timestamp
```

## Repository Options

**Repository Cipher Type Option ( `--repo-cipher-type` )** Cipher used to encrypt the repository.

The following repository types are supported:

- none - The repository is not encrypted
- aes-256-cbc - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. S3) supports encryption.

```
default: none
example: --repo1-cipher-type=aes-256-cbc
```

**Repository Host Option ( `--repo-host` )** Repository host when operating remotely via SSH.

Make sure that trusted SSH authentication is configured between the PostgreSQL host and the repository host.

When backing up and archiving to a locally mounted filesystem this setting is not required.

```
example: --repo1-host=repo1.domain.com
```

Deprecated Name: backup-host

**Repository Host Command Option ( `--repo-host-cmd` )** pgBackRest exe path on the repository host.

Required only if the path to pgbackrest is different on the local and repository hosts. If not defined, the repository host exe path will be set the same as the local exe path.

```
example: --repo1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: backup-cmd

**Repository Host Configuration Option ( `--repo-host-config` )** pgBackRest repository host configuration file.

Sets the location of the configuration file on the repository host. This is only required if the repository host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --repo1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: backup-config

**Repository Host Configuration Include Path Option ( `--repo-host-config-include-path` )** pgBackRest repository host configuration include path.

Sets the location of the configuration include path on the repository host. This is only required if the repository host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: --repo1-host-config-include-path=/conf/pgbackrest/conf.d
```

**Repository Host Configuration Path Option ( `--repo-host-config-path` )** pgBackRest repository host configuration path.

Sets the location of the configuration path on the repository host. This is only required if the repository host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --repo1-host-config-path=/conf/pgbackrest
```

**Repository Host Port Option ( `--repo-host-port` )** Repository host port when repo-host is set.

Use this option to specify a non-default port for the repository host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --repo1-host-port=25
```

Deprecated Name: backup-ssh-port

**Repository Host User Option ( `-repo-host-user` )** Repository host user when repo-host is set.

Defines the user that will be used for operations on the repository host. Preferably this is not the postgres user but rather some other user like pgbackrest . If PostgreSQL runs on the repository host the postgres user can be placed in the pgbackrest group so it has read permissions on the repository without being able to damage the contents accidentally.

```
default: pgbackrest
example: --repo1-host-user=repo-user
```

Deprecated Name: backup-user

**Repository Path Option ( `-repo-path` )** Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: --repo1-path=/backup/db/backrest
```

**S3 Repository Bucket Option ( `-repo-s3-bucket` )** S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting repo-path=/ but it is usually best to specify a prefix, such as /repo , so logs and other AWS generated content can also be stored in the bucket.

```
example: --repo1-s3-bucket=pg-backup
```

**S3 SSL CA File Option ( `-repo-s3-ca-file` )** S3 SSL CA File.

Use a CA file other than the system default.

```
example: --repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

**S3 SSL CA Path Option ( `-repo-s3-ca-path` )** S3 SSL CA Path.

Use a CA path other than the system default.

```
example: --repo1-s3-ca-path=/etc/pki/tls/certs
```

**S3 Repository Endpoint Option ( `-repo-s3-endpoint` )** S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: --repo1-s3-endpoint=s3.amazonaws.com
```

**S3 Repository Host Option ( `-repo-s3-host` )** S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: --repo1-s3-host=127.0.0.1
```

**S3 Repository Region Option ( `-repo-s3-region` )** S3 repository region.

The AWS region where the bucket was created.

```
example: --repo1-s3-region=us-east-1
```

**S3 Repository Verify SSL Option ( `--repo-s3-verify-ssl` )** Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
example: --no-repo1-s3-verify-ssl
```

**Repository Type Option ( `--repo-type` )** Type of storage used for the repository.

The following repository types are supported:

- cifs - Like posix , but disables links and directory fsyncs
- posix - Posix-compliant file systems
- s3 - AWS Simple Storage Service

```
default: posix
example: --repo1-type=cifs
```

### Stanza Options

**PostgreSQL Host Option ( `--pg-host` )** PostgreSQL host for operating remotely via SSH.

Used for backups where the PostgreSQL host is different from the repository host.

```
example: --pg1-host=db.domain.com
```

Deprecated Name: db-host

**PostgreSQL Host Command Option ( `--pg-host-cmd` )** pgBackRest exe path on the PostgreSQL host.

Required only if the path to pgbackrest is different on the local and PostgreSQL hosts. If not defined, the database host exe path will be set the same as the local exe path.

```
example: --pg1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: db-cmd

**PostgreSQL Host Configuration Option ( `--pg-host-config` )** pgBackRest database host configuration file.

Sets the location of the configuration file on the PostgreSQL host. This is only required if the PostgreSQL host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --pg1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: db-config

**PostgreSQL Host Configuration Include Path Option ( `--pg-host-config-include-path` )** pgBackRest database host configuration include path.

Sets the location of the configuration include path on the PostgreSQL host. This is only required if the PostgreSQL host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: --pg1-host-config-include-path=/conf/pgbackrest/conf.d
```

**PostgreSQL Host Configuration Path Option ( `--pg-host-config-path` )** pgBackRest database host configuration path.

Sets the location of the configuration path on the PostgreSQL host. This is only required if the PostgreSQL host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --pg1-host-config-path=/conf/pgbackrest
```

**PostgreSQL Host Port Option ( `-pg-host-port` )** PostgreSQL host port when pg-host is set.

Use this option to specify a non-default port for the PostgreSQL host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --pg1-host-port=25
```

Deprecated Name: db-ssh-port

**PostgreSQL Host User Option ( `-pg-host-user` )** PostgreSQL host logon user when pg-host is set.

This user will also own the remote pgBackRest process and will initiate connections to PostgreSQL . For this to work correctly the user should be the PostgreSQL database cluster owner which is generally postgres , the default.

```
default: postgres
example: --pg1-host-user=db_owner
```

Deprecated Name: db-user

**PostgreSQL Path Option ( `-pg-path` )** PostgreSQL data directory.

This should be the same as the data\_directory setting in postgresql.conf . Even though this value can be read from postgresql.conf or PostgreSQL it is prudent to set it in case those resources are not available during a restore or offline backup scenario.

The pg-path option is tested against the value reported by PostgreSQL on every online backup so it should always be current.

```
example: --pg1-path=/data/db
```

Deprecated Name: db-path

**PostgreSQL Port Option ( `-pg-port` )** PostgreSQL port.

Port that PostgreSQL is running on. This usually does not need to be specified as most PostgreSQL clusters run on the default port.

```
default: 5432
allowed: 0-65535
example: --pg1-port=6543
```

Deprecated Name: db-port

**PostgreSQL Socket Path Option ( `-pg-socket-path` )** PostgreSQL unix socket path.

The unix socket directory that was specified when PostgreSQL was started. pgBackRest will automatically look in the standard location for your OS so there is usually no need to specify this setting unless the socket directory was explicitly modified with the unix\_socket\_directory setting in postgresql.conf .

```
allowed: 0-65535
example: --pg1-socket-path=/var/run/postgresql
```

Deprecated Name: db-socket-path

## Expire Command ( `expire` )

pgBackRest does backup rotation but is not concerned with when the backups were created. If two full backups are configured for retention, pgBackRest will keep two full backups no matter whether they occur two hours or two weeks apart.

## General Options

**Buffer Size Option ( `-buffer-size` )** Buffer size for file operations.

Set the buffer size used for copy, compress, and uncompress functions. A maximum of 3 buffers will be in use at a time per process. An additional maximum of 256K per process may be used for zlib buffers.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024. For example, the case-insensitive value 32k (or 32KB) can be used instead of 32768.

Allowed values, in bytes, are 16384 , 32768 , 65536 , 131072 , 262144 , 524288 , 1048576 , 2097152 , 4194304 , 8388608 , and 16777216 .



```
default: 4194304
example: --buffer-size=32K
```

**SSH client command Option ( `-cmd-ssh` )** Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the ssh executable is not in \$PATH.

```
default: ssh
example: --cmd-ssh=/usr/bin/ssh
```

**Config Option ( `-config` )** pgBackRest configuration file.

Use this option to specify a different configuration file than the default.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --config=/conf/pgbackrest/pgbackrest.conf
```

**Config Include Path Option ( `-config-include-path` )** Path to additional pgBackRest configuration files.

Configuration files existing in the specified location with extension .conf will be concatenated with the pgBackRest configuration file, resulting in one configuration file.

```
default: /etc/pgbackrest/conf.d
example: --config-include-path=/conf/pgbackrest/conf.d
```

**Config Path Option ( `-config-path` )** Base path of pgBackRest configuration files.

This setting is used to override the default base path setting for the `-config` and `-config-include-path` options unless they are explicitly set on the command-line.

For example, passing only `-config-path=/conf/pgbackrest` results in the `-config` default being set to `/conf/pgbackrest/pgbackrest.conf` and the `-config-include-path` default being set to `/conf/pgbackrest/conf.d`.

```
default: /etc/pgbackrest
example: --config-path=/conf/pgbackrest
```

**Lock Path Option ( `-lock-path` )** Path where lock files are stored.

The lock path provides a location for pgBackRest to create lock files to prevent conflicting operations from being run concurrently.

```
default: /tmp/pgbackrest
example: --lock-path=/backup/db/lock
```

**Neutral Umask Option ( `-neutral-umask` )** Use a neutral umask.

Sets the umask to 0000 so modes in the repository are created in a sensible way. The default directory mode is 0750 and default file mode is 0640. The lock and log directories set the directory and file mode to 0770 and 0660 respectively.

To use the executing user's umask instead specify `neutral-umask=n` in the config file or `-no-neutral-umask` on the command line.

```
default: y
example: --no-neutral-umask
```

**Stanza Option ( `-stanza` )** Defines the stanza.

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as `app` or `dw`, rather than the local cluster name, such as `main` or `prod`.

```
example: --stanza=main
```

## Log Options

**Console Log Level Option ( `-log-level-console` )** Level for console logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-console=error
```

**File Log Level Option ( `-log-level-file` )** Level for file logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
```

```
example: --log-level-file=debug
```

**Std Error Log Level Option ( `-log-level-stderr` )** Level for stderr logging.

Specifies which log levels will output to stderr rather than stdout (specified by `log-level-console`). The timestamp and process will not be output to stderr.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-stderr=error
```

**Log Path Option ( `-log-path` )** Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
```

```
example: --log-path=/backup/db/log
```

**Log Subprocesses Option ( `-log-subprocess` )** Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file` .

```
default: n
example: --log-subprocess
```

**Log Timestamp Option ( `-log-timestamp` )** Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
example: --no-log-timestamp
```

## Repository Options

**Repository Cipher Type Option ( `-repo-cipher-type` )** Cipher used to encrypt the repository.

The following repository types are supported:

- none - The repository is not encrypted
- aes-256-cbc - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. S3) supports encryption.

```
default: none
example: --repo1-cipher-type=aes-256-cbc
```

**Repository Path Option ( `-repo-path` )** Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: --repo1-path=/backup/db/backrest
```

**Archive Retention Option ( `-repo-retention-archive` )** Number of backups worth of continuous WAL to retain.

NOTE:

WAL segments required to make a backup consistent are always retained until the backup is expired regardless of how this option is configured.

If this value is not set, then the archive to expire will default to the `repo-retention-full` (or `repo-retention-diff`) value corresponding to the `repo-retention-archive-type` if set to full (or diff). This will ensure that WAL is only expired for backups that are already expired.

This option must be set if `repo-retention-archive-type` is set to `incr` . If disk space is at a premium, then this setting, in conjunction with `repo-retention-archive-type` , can be used to aggressively expire WAL segments. However, doing so negates the ability to perform PITR from the backups with expired WAL and is therefore **not** recommended.

```
allowed: 1-9999999
example: --repo1-retention-archive=2
```

Deprecated Name: `retention-archive`

**Archive Retention Type Option ( `--repo-retention-archive-type` )** Backup type for WAL retention.

If set to full pgBackRest will keep archive logs for the number of full backups defined by `repo-retention-archive` . If set to diff (differential) pgBackRest will keep archive logs for the number of full and differential backups defined by `repo-retention-archive` , meaning if the last backup taken was a full backup, it will be counted as a differential for the purpose of `repo-retention`. If set to incr (incremental) pgBackRest will keep archive logs for the number of full, differential, and incremental backups defined by `repo-retention-archive` . It is recommended that this setting not be changed from the default which will only expire WAL in conjunction with expiring full backups.

```
default: full
example: --repo1-retention-archive-type=diff
```

Deprecated Name: `retention-archive-type`

**Differential Retention Option ( `--repo-retention-diff` )** Number of differential backups to retain.

When a differential backup expires, all incremental backups associated with the differential backup will also expire. When not defined all differential backups will be kept until the full backups they depend on expire.

```
allowed: 1-9999999
example: --repo1-retention-diff=3
```

Deprecated Name: `retention-diff`

**Full Retention Option ( `--repo-retention-full` )** Number of full backups to retain.

When a full backup expires, all differential and incremental backups associated with the full backup will also expire. When the option is not defined a warning will be issued. If indefinite retention is desired then set the option to the max value.

```
allowed: 1-9999999
example: --repo1-retention-full=2
```

Deprecated Name: `retention-full`

**S3 Repository Bucket Option ( `--repo-s3-bucket` )** S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting `repo-path=/` but it is usually best to specify a prefix, such as `/repo` , so logs and other AWS generated content can also be stored in the bucket.

```
example: --repo1-s3-bucket=pg-backup
```

**S3 SSL CA File Option ( `--repo-s3-ca-file` )** S3 SSL CA File.

Use a CA file other than the system default.

```
example: --repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

**S3 SSL CA Path Option ( `--repo-s3-ca-path` )** S3 SSL CA Path.

Use a CA path other than the system default.

```
example: --repo1-s3-ca-path=/etc/pki/tls/certs
```

**S3 Repository Endpoint Option ( `--repo-s3-endpoint` )** S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: --repo1-s3-endpoint=s3.amazonaws.com
```

**S3 Repository Host Option ( `--repo-s3-host` )** S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: --repo1-s3-host=127.0.0.1
```

**S3 Repository Region Option ( `-repo-s3-region` )** S3 repository region.

The AWS region where the bucket was created.

```
example: --repo1-s3-region=us-east-1
```

**S3 Repository Verify SSL Option ( `-repo-s3-verify-ssl` )** Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
```

```
example: --no-repo1-s3-verify-ssl
```

**Repository Type Option ( `-repo-type` )** Type of storage used for the repository.

The following repository types are supported:

- cifs - Like posix , but disables links and directory fsyncs
- posix - Posix-compliant file systems
- s3 - AWS Simple Storage Service

```
default: posix
```

```
example: --repo1-type=cifs
```

## Stanza Options

**PostgreSQL Host Option ( `-pg-host` )** PostgreSQL host for operating remotely via SSH.

Used for backups where the PostgreSQL host is different from the repository host.

```
example: --pg1-host=db.domain.com
```

Deprecated Name: db-host

**PostgreSQL Host Command Option ( `-pg-host-cmd` )** pgBackRest exe path on the PostgreSQL host.

Required only if the path to pgbackrest is different on the local and PostgreSQL hosts. If not defined, the database host exe path will be set the same as the local exe path.

```
example: --pg1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: db-cmd

**PostgreSQL Host Configuration Option ( `-pg-host-config` )** pgBackRest database host configuration file.

Sets the location of the configuration file on the PostgreSQL host. This is only required if the PostgreSQL host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
```

```
example: --pg1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: db-config

**PostgreSQL Host Configuration Include Path Option ( `-pg-host-config-include-path` )** pgBackRest database host configuration include path.

Sets the location of the configuration include path on the PostgreSQL host. This is only required if the PostgreSQL host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
```

```
example: --pg1-host-config-include-path=/conf/pgbackrest/conf.d
```

**PostgreSQL Host Configuration Path Option ( `-pg-host-config-path` )** pgBackRest database host configuration path.

Sets the location of the configuration path on the PostgreSQL host. This is only required if the PostgreSQL host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --pg1-host-config-path=/conf/pgbackrest
```

**PostgreSQL Host Port Option ( `-pg-host-port` )** PostgreSQL host port when pg-host is set.

Use this option to specify a non-default port for the PostgreSQL host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --pg1-host-port=25
```

Deprecated Name: db-ssh-port

**PostgreSQL Host User Option ( `-pg-host-user` )** PostgreSQL host logon user when pg-host is set.

This user will also own the remote pgBackRest process and will initiate connections to PostgreSQL . For this to work correctly the user should be the PostgreSQL database cluster owner which is generally postgres , the default.

```
default: postgres
example: --pg1-host-user=db_owner
```

Deprecated Name: db-user

## Help Command ( `help` )

Three levels of help are provided. If no command is specified then general help will be displayed. If a command is specified then a full description of the command will be displayed along with a list of valid options. If an option is specified in addition to a command then the a full description of the option as it applies to the command will be displayed.

## Info Command ( `info` )

The info command operates on a single stanza or all stanzas. Text output is the default and gives a human-readable summary of backups for the stanza(s) requested. This format is subject to change with any release.

For machine-readable output use `-output=json` . The JSON output contains far more information than the text output and is kept stable unless a bug is found.

## Command Options

**Output Option ( `-output` )** Output format.

The following output types are supported:

- text - Human-readable summary of backup information.
- json - Exhaustive machine-readable backup information in JSON format.

```
default: text
example: --output=json
```

## General Options

**Buffer Size Option ( `-buffer-size` )** Buffer size for file operations.

Set the buffer size used for copy, compress, and uncompress functions. A maximum of 3 buffers will be in use at a time per process. An additional maximum of 256K per process may be used for zlib buffers.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024. For example, the case-insensitive value 32k (or 32KB) can be used instead of 32768.

Allowed values, in bytes, are 16384 , 32768 , 65536 , 131072 , 262144 , 524288 , 1048576 , 2097152 , 4194304 , 8388608 , and 16777216 .

```
default: 4194304
example: --buffer-size=32K
```

**SSH client command Option ( `-cmd-ssh` )** Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the ssh executable is not in \$PATH.

```
default: ssh
example: --cmd-ssh=/usr/bin/ssh
```

**Compress Level Option ( `-compress-level` )** Compression level for stored files.

Sets the zlib level to be used for file compression when `compress=y` .

```
default: 6
allowed: 0-9
example: --compress-level=9
```

**Network Compress Level Option ( `-compress-level-network` )** Compression level for network transfer when `compress=n` .

Sets the zlib level to be used for protocol compression when `compress=n` and the database cluster is not on the same host as the repository. Protocol compression is used to reduce network traffic but can be disabled by setting `compress-level-network=0` . When `compress=y` the `compress-level-network` setting is ignored and `compress-level` is used instead so that the file is only compressed once. SSH compression is always disabled.

```
default: 3
allowed: 0-9
example: --compress-level-network=1
```

**Config Option ( `-config` )** pgBackRest configuration file.

Use this option to specify a different configuration file than the default.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --config=/conf/pgbackrest/pgbackrest.conf
```

**Config Include Path Option ( `-config-include-path` )** Path to additional pgBackRest configuration files.

Configuration files existing in the specified location with extension `.conf` will be concatenated with the pgBackRest configuration file, resulting in one configuration file.

```
default: /etc/pgbackrest/conf.d
example: --config-include-path=/conf/pgbackrest/conf.d
```

**Config Path Option ( `-config-path` )** Base path of pgBackRest configuration files.

This setting is used to override the default base path setting for the `-config` and `-config-include-path` options unless they are explicitly set on the command-line.

For example, passing only `-config-path=/conf/pgbackrest` results in the `-config` default being set to `/conf/pgbackrest/pgbackrest.conf` and the `-config-include-path` default being set to `/conf/pgbackrest/conf.d` .

```
default: /etc/pgbackrest
example: --config-path=/conf/pgbackrest
```

**Lock Path Option ( `-lock-path` )** Path where lock files are stored.

The lock path provides a location for pgBackRest to create lock files to prevent conflicting operations from being run concurrently.

```
default: /tmp/pgbackrest
example: --lock-path=/backup/db/lock
```

**Protocol Timeout Option ( `-protocol-timeout` )** Protocol timeout.

Sets the timeout, in seconds, that the local or remote process will wait for a new message to be received on the protocol layer. This prevents processes from waiting indefinitely for a message. The `protocol-timeout` option must be greater than the `db-timeout` option.

```
default: 1830
allowed: 0.1-604800
example: --protocol-timeout=630
```

**Stanza Option ( `-stanza` )** Defines the stanza.

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as `app` or `dw`, rather than the local cluster name, such as `main` or `prod`.

```
example: --stanza=main
```

## Log Options

**Console Log Level Option ( `-log-level-console` )** Level for console logging.

The following log levels are supported:

- `off` - No logging at all (not recommended)
- `error` - Log only errors
- `warn` - Log warnings and errors
- `info` - Log info, warnings, and errors
- `detail` - Log detail, info, warnings, and errors
- `debug` - Log debug, detail, info, warnings, and errors
- `trace` - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
example: --log-level-console=error
```

**File Log Level Option ( `-log-level-file` )** Level for file logging.

The following log levels are supported:

- `off` - No logging at all (not recommended)
- `error` - Log only errors
- `warn` - Log warnings and errors
- `info` - Log info, warnings, and errors
- `detail` - Log detail, info, warnings, and errors
- `debug` - Log debug, detail, info, warnings, and errors
- `trace` - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
example: --log-level-file=debug
```

**Std Error Log Level Option ( `-log-level-stderr` )** Level for stderr logging.

Specifies which log levels will output to `stderr` rather than `stdout` (specified by `log-level-console`). The timestamp and process will not be output to `stderr`.

The following log levels are supported:



- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
example: --log-level-stderr=error
```

**Log Path Option ( `-log-path` )** Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
example: --log-path=/backup/db/log
```

**Log Subprocesses Option ( `-log-subprocess` )** Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file` .

```
default: n
example: --log-subprocess
```

**Log Timestamp Option ( `-log-timestamp` )** Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
example: --no-log-timestamp
```

## Repository Options

**Repository Cipher Type Option ( `-repo-cipher-type` )** Cipher used to encrypt the repository.

The following repository types are supported:

- none - The repository is not encrypted
- aes-256-cbc - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. S3) supports encryption.

```
default: none
example: --repo1-cipher-type=aes-256-cbc
```

**Repository Host Option ( `-repo-host` )** Repository host when operating remotely via SSH.

Make sure that trusted SSH authentication is configured between the PostgreSQL host and the repository host.

When backing up and archiving to a locally mounted filesystem this setting is not required.

```
example: --repo1-host=repo1.domain.com
```

Deprecated Name: `backup-host`

**Repository Host Command Option ( `-repo-host-cmd` )** pgBackRest exe path on the repository host.

Required only if the path to pgbackrest is different on the local and repository hosts. If not defined, the repository host exe path will be set the same as the local exe path.

```
example: --repo1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: `backup-cmd`

**Repository Host Configuration Option ( `-repo-host-config` )** pgBackRest repository host configuration file.

Sets the location of the configuration file on the repository host. This is only required if the repository host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --repo1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: backup-config

**Repository Host Configuration Include Path Option ( `-repo-host-config-include-path` )** pgBackRest repository host configuration include path.

Sets the location of the configuration include path on the repository host. This is only required if the repository host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: --repo1-host-config-include-path=/conf/pgbackrest/conf.d
```

**Repository Host Configuration Path Option ( `-repo-host-config-path` )** pgBackRest repository host configuration path.

Sets the location of the configuration path on the repository host. This is only required if the repository host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --repo1-host-config-path=/conf/pgbackrest
```

**Repository Host Port Option ( `-repo-host-port` )** Repository host port when repo-host is set.

Use this option to specify a non-default port for the repository host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --repo1-host-port=25
```

Deprecated Name: backup-ssh-port

**Repository Host User Option ( `-repo-host-user` )** Repository host user when repo-host is set.

Defines the user that will be used for operations on the repository host. Preferably this is not the postgres user but rather some other user like pgbackrest . If PostgreSQL runs on the repository host the postgres user can be placed in the pgbackrest group so it has read permissions on the repository without being able to damage the contents accidentally.

```
default: pgbackrest
example: --repo1-host-user=repo-user
```

Deprecated Name: backup-user

**Repository Path Option ( `-repo-path` )** Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: --repo1-path=/backup/db/backrest
```

**S3 Repository Bucket Option ( `-repo-s3-bucket` )** S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting repo-path=/ but it is usually best to specify a prefix, such as /repo , so logs and other AWS generated content can also be stored in the bucket.

```
example: --repo1-s3-bucket=pg-backup
```

### **S3 SSL CA File Option ( `--repo-s3-ca-file` )** S3 SSL CA File.

Use a CA file other than the system default.

```
example: --repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

### **S3 SSL CA Path Option ( `--repo-s3-ca-path` )** S3 SSL CA Path.

Use a CA path other than the system default.

```
example: --repo1-s3-ca-path=/etc/pki/tls/certs
```

### **S3 Repository Endpoint Option ( `--repo-s3-endpoint` )** S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: --repo1-s3-endpoint=s3.amazonaws.com
```

### **S3 Repository Host Option ( `--repo-s3-host` )** S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: --repo1-s3-host=127.0.0.1
```

### **S3 Repository Region Option ( `--repo-s3-region` )** S3 repository region.

The AWS region where the bucket was created.

```
example: --repo1-s3-region=us-east-1
```

### **S3 Repository Verify SSL Option ( `--repo-s3-verify-ssl` )** Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
```

```
example: --no-repo1-s3-verify-ssl
```

### **Repository Type Option ( `--repo-type` )** Type of storage used for the repository.

The following repository types are supported:

- cifs - Like posix , but disables links and directory fsyncs
- posix - Posix-compliant file systems
- s3 - AWS Simple Storage Service

```
default: posix
```

```
example: --repo1-type=cifs
```

## **Restore Command ( `restore` )**

This command is generally run manually, but there are instances where it might be automated.

### **Command Options**

#### **Include Database Option ( `--db-include` )** Restore only specified databases.

This feature allows only selected databases to be restored. Databases not specifically included will be restored as sparse, zeroed files to save space but still allow PostgreSQL to perform recovery. After recovery the databases that were not included will not be accessible but can be removed with the drop database command.

NOTE:

built-in databases ( `template0` , `template1` , and `postgres` ) are always restored.

The `--db-include` option can be passed multiple times to specify more than one database to include.

```
example: --db-include=db_main
```

**Force Option ( `-force` )** Force a restore.

By itself this option forces the PostgreSQL data and tablespace paths to be completely overwritten. In combination with `-delta` a timestamp/size delta will be performed instead of using checksums.

```
default: n
example: --force
```

**Link All Option ( `-link-all` )** Restore all symlinks.

By default symlinked directories and files are restored as normal directories and files in `$PGDATA`. This is because it may not be safe to restore symlinks to their original destinations on a system other than where the original backup was performed. This option restores all the symlinks just as they were on the original system where the backup was performed.

```
default: n
example: --link-all
```

**Link Map Option ( `-link-map` )** Modify the destination of a symlink.

Allows the destination file or path of a symlink to be changed on restore. This is useful for restoring to systems that have a different storage layout than the original system where the backup was generated.

```
example: --link-map=pg_xlog=/data/xlog
```

**Recovery Option Option ( `-recovery-option` )** Set an option in `recovery.conf` .

See <http://www.postgresql.org/docs/X.X/static/recovery-config.html> for details on `recovery.conf` options (replace X.X with your PostgreSQL version). This option can be used multiple times.

NOTE:

The `restore_command` option will be automatically generated but can be overridden with this option. Be careful about specifying your own `restore_command` as `pgBackRest` is designed to handle this for you. Target Recovery options (`recovery_target_name`, `recovery_target_time`, etc.) are generated automatically by `pgBackRest` and should not be set with this option.

Since `pgBackRest` does not start PostgreSQL after writing the `recovery.conf` file, it is always possible to edit/check `recovery.conf` before manually restarting.

```
example: --recovery-option=primary_conninfo=db.mydomain.com
```

**Set Option ( `-set` )** Backup set to restore.

The backup set to be restored. `latest` will restore the latest backup, otherwise provide the name of the backup to restore.

```
default: latest
example: --set=20150131-153358F_20150131-153401I
```

**Tablespace Map Option ( `-tablespace-map` )** Restore a tablespace into the specified directory.

Moves a tablespace to a new location during the restore. This is useful when tablespace locations are not the same on a replica, or an upgraded system has different mount points.

Since PostgreSQL 9.2 tablespace locations are not stored in `pg_tablespace` so moving tablespaces can be done with impunity. However, moving a tablespace to the `data_directory` is not recommended and may cause problems. For more information on moving tablespaces <http://www.databasesoup.com/2013/11/moving-tablespaces.html> is a good resource.

```
example: --tablespace-map=ts_01=/db/ts_01
```

**Map All Tablespaces Option ( `-tablespace-map-all` )** Restore all tablespaces into the specified directory.

By default tablespaces are restored into their original locations and while this behavior can be modified by with the `tablespace-map` open it is sometime preferable to remap all tablespaces to a new directory all at once. This is particularly useful for development or staging systems that may not have the same storage layout as the original system where the backup was generated.

The path specified will be the parent path used to create all the tablespaces in the backup.

```
example: --tablespace-map-all=/data/tablespace
```

**Target Option ( `-target` )** Recovery target.

Defines the recovery target when `-type` is `name` , `xid` , or `time` .

```
example: --target=2015-01-30 14:15:11 EST
```

**Target Action Option ( `-target-action` )** Action to take when recovery target is reached.

This option is effective when `hot_standby=on` is configured in `postgresql.conf` , otherwise the cluster will be promoted when the target is reached or there is no more WAL in the archive.

The following actions are supported:

- `pause` - pause when recovery target is reached. ( PostgreSQL  $\geq$  9.1)
- `promote` - promote and switch timeline when recovery target is reached. ( PostgreSQL  $\geq$  9.1)
- `shutdown` - shutdown server when recovery target is reached. ( PostgreSQL  $\geq$  9.5)

```
default: pause
```

```
example: --target-action=promote
```

**Target Exclusive Option ( `-target-exclusive` )** Stop just before the recovery target is reached.

Defines whether recovery to the target would be exclusive (the default is inclusive) and is only valid when `-type` is `time` or `xid` . For example, using `-target-exclusive` would exclude the contents of transaction 1007 when `-type=xid` and `-target=1007` . See the `recovery_target_inclusive` option in the PostgreSQL docs for more information.

```
default: n
```

```
example: --no-target-exclusive
```

**Target Timeline Option ( `-target-timeline` )** Recover along a timeline.

See `recovery_target_timeline` in the PostgreSQL docs for more information.

```
example: --target-timeline=3
```

**Type Option ( `-type` )** Recovery type.

The following recovery types are supported:

- `default` - recover to the end of the archive stream.
- `immediate` - recover only until the database becomes consistent. This option is only supported on PostgreSQL  $\geq$  9.4.
- `name` - recover the restore point specified in `-target` .
- `xid` - recover to the transaction id specified in `-target` .
- `time` - recover to the time specified in `-target` .
- `preserve` - preserve the existing `recovery.conf` file.
- `none` - no `recovery.conf` file is written so PostgreSQL will attempt to achieve consistency using WAL segments present in `pg_xlog` / `pg_wal` . Provide the required WAL segments or use the `archive-copy` setting to include them with the backup.

```
default: default
```

```
example: --type=xid
```

## General Options

**Buffer Size Option ( `-buffer-size` )** Buffer size for file operations.

Set the buffer size used for copy, compress, and uncompress functions. A maximum of 3 buffers will be in use at a time per process. An additional maximum of 256K per process may be used for zlib buffers.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024. For example, the case-insensitive value `32k` (or `32KB`) can be used instead of `32768`.

Allowed values, in bytes, are 16384 , 32768 , 65536 , 131072 , 262144 , 524288 , 1048576 , 2097152 , 4194304 , 8388608 , and 16777216 .

```
default: 4194304
example: --buffer-size=32K
```

**SSH client command Option ( `-cmd-ssh` )** Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the ssh executable is not in \$PATH.

```
default: ssh
example: --cmd-ssh=/usr/bin/ssh
```

**Compress Option ( `-compress` )** Use gzip file compression.

Backup files are compatible with command-line gzip tools.

```
default: y
example: --no-compress
```

**Compress Level Option ( `-compress-level` )** Compression level for stored files.

Sets the zlib level to be used for file compression when `compress=y` .

```
default: 6
allowed: 0-9
example: --compress-level=9
```

**Network Compress Level Option ( `-compress-level-network` )** Compression level for network transfer when `compress=n` .

Sets the zlib level to be used for protocol compression when `compress=n` and the database cluster is not on the same host as the repository. Protocol compression is used to reduce network traffic but can be disabled by setting `compress-level-network=0` . When `compress=y` the `compress-level-network` setting is ignored and `compress-level` is used instead so that the file is only compressed once. SSH compression is always disabled.

```
default: 3
allowed: 0-9
example: --compress-level-network=1
```

**Config Option ( `-config` )** pgBackRest configuration file.

Use this option to specify a different configuration file than the default.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --config=/conf/pgbackrest/pgbackrest.conf
```

**Config Include Path Option ( `-config-include-path` )** Path to additional pgBackRest configuration files.

Configuration files existing in the specified location with extension `.conf` will be concatenated with the pgBackRest configuration file, resulting in one configuration file.

```
default: /etc/pgbackrest/conf.d
example: --config-include-path=/conf/pgbackrest/conf.d
```

**Config Path Option ( `-config-path` )** Base path of pgBackRest configuration files.

This setting is used to override the default base path setting for the `-config` and `-config-include-path` options unless they are explicitly set on the command-line.

For example, passing only `-config-path=/conf/pgbackrest` results in the `-config` default being set to `/conf/pgbackrest/pgbackrest.conf` and the `-config-include-path` default being set to `/conf/pgbackrest/conf.d` .

```
default: /etc/pgbackrest
example: --config-path=/conf/pgbackrest
```

**Delta Option ( `-delta` )** Restore or backup using checksums.

During a restore, by default the PostgreSQL data and tablespace directories are expected to be present but empty. This option performs a delta restore using checksums.

During a backup, this option will use checksums instead of the timestamps to determine if files will be copied.

```
default: n
example: --delta
```

**Lock Path Option ( `-lock-path` )** Path where lock files are stored.

The lock path provides a location for pgBackRest to create lock files to prevent conflicting operations from being run concurrently.

```
default: /tmp/pgbackrest
example: --lock-path=/backup/db/lock
```

**Neutral Umask Option ( `-neutral-umask` )** Use a neutral umask.

Sets the umask to 0000 so modes in the repository are created in a sensible way. The default directory mode is 0750 and default file mode is 0640. The lock and log directories set the directory and file mode to 0770 and 0660 respectively.

To use the executing user's umask instead specify `neutral-umask=n` in the config file or `-no-neutral-umask` on the command line.

```
default: y
example: --no-neutral-umask
```

**Process Maximum Option ( `-process-max` )** Max processes to use for compress/transfer.

Each process will perform compression and transfer to make the command run faster, but don't set `process-max` so high that it impacts database performance.

```
default: 1
allowed: 1-96
example: --process-max=4
```

**Protocol Timeout Option ( `-protocol-timeout` )** Protocol timeout.

Sets the timeout, in seconds, that the local or remote process will wait for a new message to be received on the protocol layer. This prevents processes from waiting indefinitely for a message. The `protocol-timeout` option must be greater than the `db-timeout` option.

```
default: 1830
allowed: 0.1-604800
example: --protocol-timeout=630
```

**Stanza Option ( `-stanza` )** Defines the stanza.

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as `app` or `dw`, rather than the local cluster name, such as `main` or `prod`.

```
example: --stanza=main
```

## Log Options

**Console Log Level Option ( `-log-level-console` )** Level for console logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-console=error
```

**File Log Level Option ( `-log-level-file` )** Level for file logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
```

```
example: --log-level-file=debug
```

**Std Error Log Level Option ( `-log-level-stderr` )** Level for stderr logging.

Specifies which log levels will output to stderr rather than stdout (specified by `log-level-console`). The timestamp and process will not be output to stderr.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-stderr=error
```

**Log Path Option ( `-log-path` )** Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
```

```
example: --log-path=/backup/db/log
```



**Log Subprocesses Option ( `-log-subprocess` )** Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file` .

```
default: n
example: --log-subprocess
```

**Log Timestamp Option ( `-log-timestamp` )** Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
example: --no-log-timestamp
```

## Repository Options

**Repository Cipher Type Option ( `-repo-cipher-type` )** Cipher used to encrypt the repository.

The following repository types are supported:

- none - The repository is not encrypted
- aes-256-cbc - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. S3) supports encryption.

```
default: none
example: --repo1-cipher-type=aes-256-cbc
```

**Repository Host Option ( `-repo-host` )** Repository host when operating remotely via SSH.

Make sure that trusted SSH authentication is configured between the PostgreSQL host and the repository host.

When backing up and archiving to a locally mounted filesystem this setting is not required.

```
example: --repo1-host=repo1.domain.com
```

Deprecated Name: `backup-host`

**Repository Host Command Option ( `-repo-host-cmd` )** `pgBackRest` exe path on the repository host.

Required only if the path to `pgbackrest` is different on the local and repository hosts. If not defined, the repository host exe path will be set the same as the local exe path.

```
example: --repo1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: `backup-cmd`

**Repository Host Configuration Option ( `-repo-host-config` )** `pgBackRest` repository host configuration file.

Sets the location of the configuration file on the repository host. This is only required if the repository host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --repo1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: `backup-config`

**Repository Host Configuration Include Path Option ( `-repo-host-config-include-path` )** `pgBackRest` repository host configuration include path.

Sets the location of the configuration include path on the repository host. This is only required if the repository host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: --repo1-host-config-include-path=/conf/pgbackrest/conf.d
```

**Repository Host Configuration Path Option ( `--repo-host-config-path` )** pgBackRest repository host configuration path.

Sets the location of the configuration path on the repository host. This is only required if the repository host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --repo1-host-config-path=/conf/pgbackrest
```

**Repository Host Port Option ( `--repo-host-port` )** Repository host port when `repo-host` is set.

Use this option to specify a non-default port for the repository host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --repo1-host-port=25
```

Deprecated Name: `backup-ssh-port`

**Repository Host User Option ( `--repo-host-user` )** Repository host user when `repo-host` is set.

Defines the user that will be used for operations on the repository host. Preferably this is not the postgres user but rather some other user like `pgbackrest`. If PostgreSQL runs on the repository host the postgres user can be placed in the `pgbackrest` group so it has read permissions on the repository without being able to damage the contents accidentally.

```
default: pgbackrest
example: --repo1-host-user=repo-user
```

Deprecated Name: `backup-user`

**Repository Path Option ( `--repo-path` )** Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: --repo1-path=/backup/db/backrest
```

**S3 Repository Bucket Option ( `--repo-s3-bucket` )** S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting `repo-path=` but it is usually best to specify a prefix, such as `/repo`, so logs and other AWS generated content can also be stored in the bucket.

```
example: --repo1-s3-bucket=pg-backup
```

**S3 SSL CA File Option ( `--repo-s3-ca-file` )** S3 SSL CA File.

Use a CA file other than the system default.

```
example: --repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

**S3 SSL CA Path Option ( `--repo-s3-ca-path` )** S3 SSL CA Path.

Use a CA path other than the system default.

```
example: --repo1-s3-ca-path=/etc/pki/tls/certs
```

**S3 Repository Endpoint Option ( `--repo-s3-endpoint` )** S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: --repo1-s3-endpoint=s3.amazonaws.com
```

**S3 Repository Host Option ( `-repo-s3-host` )** S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: --repo1-s3-host=127.0.0.1
```

**S3 Repository Region Option ( `-repo-s3-region` )** S3 repository region.

The AWS region where the bucket was created.

```
example: --repo1-s3-region=us-east-1
```

**S3 Repository Verify SSL Option ( `-repo-s3-verify-ssl` )** Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
```

```
example: --no-repo1-s3-verify-ssl
```

**Repository Type Option ( `-repo-type` )** Type of storage used for the repository.

The following repository types are supported:

- cifs - Like posix , but disables links and directory fsyncs
- posix - Posix-compliant file systems
- s3 - AWS Simple Storage Service

```
default: posix
```

```
example: --repo1-type=cifs
```

## Stanza Options

**PostgreSQL Path Option ( `-pg-path` )** PostgreSQL data directory.

This should be the same as the `data_directory` setting in `postgresql.conf` . Even though this value can be read from `postgresql.conf` or PostgreSQL it is prudent to set it in case those resources are not available during a restore or offline backup scenario.

The `pg-path` option is tested against the value reported by PostgreSQL on every online backup so it should always be current.

```
example: --pg1-path=/data/db
```

Deprecated Name: `db-path`

## Stanza Create Command ( `stanza-create` )

The `stanza-create` command must be run on the host where the repository is located after the stanza has been configured in `pgbackrest.conf` .

## Command Options

**Backup from Standby Option ( `-backup-standby` )** Backup from the standby cluster.

Enable backup from standby to reduce load on the primary cluster. This option requires that both the primary and standby hosts be configured.

```
default: n
```

```
example: --backup-standby
```

**Force Option ( `-force` )** Force stanza creation.

**CAUTION:**

Use `-force` only as a last resort, when all else fails. If data is missing from the repository then the recreated `.info` files will likely be corrupt.

If the required stanza `.info` files do not exist in the repository but backups or WAL segments do exist, then this option can be used to force the stanza to be created from the existing data in the repository. This is most likely to be useful after corruption or an incomplete restore of the repository from elsewhere.

```
default: n
example: --no-force
```

**Online Option ( `-online` )** Create on an online cluster.

Specifying `-no-online` prevents `pgBackRest` from connecting to PostgreSQL when creating the stanza.

```
default: y
example: --no-online
```

## General Options

**Buffer Size Option ( `-buffer-size` )** Buffer size for file operations.

Set the buffer size used for copy, compress, and uncompress functions. A maximum of 3 buffers will be in use at a time per process. An additional maximum of 256K per process may be used for zlib buffers.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024. For example, the case-insensitive value `32k` (or `32KB`) can be used instead of `32768`.

Allowed values, in bytes, are 16384 , 32768 , 65536 , 131072 , 262144 , 524288 , 1048576 , 2097152 , 4194304 , 8388608 , and 16777216 .

```
default: 4194304
example: --buffer-size=32K
```

**SSH client command Option ( `-cmd-ssh` )** Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the `ssh` executable is not in `$PATH`.

```
default: ssh
example: --cmd-ssh=/usr/bin/ssh
```

**Compress Level Option ( `-compress-level` )** Compression level for stored files.

Sets the zlib level to be used for file compression when `compress=y` .

```
default: 6
allowed: 0-9
example: --compress-level=9
```

**Network Compress Level Option ( `-compress-level-network` )** Compression level for network transfer when `compress=n` .

Sets the zlib level to be used for protocol compression when `compress=n` and the database cluster is not on the same host as the repository. Protocol compression is used to reduce network traffic but can be disabled by setting `compress-level-network=0` . When `compress=y` the `compress-level-network` setting is ignored and `compress-level` is used instead so that the file is only compressed once. SSH compression is always disabled.

```
default: 3
allowed: 0-9
example: --compress-level-network=1
```

**Config Option ( `-config` )** `pgBackRest` configuration file.

Use this option to specify a different configuration file than the default.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --config=/conf/pgbackrest/pgbackrest.conf
```

**Config Include Path Option ( `-config-include-path` )** Path to additional pgBackRest configuration files.

Configuration files existing in the specified location with extension `.conf` will be concatenated with the pgBackRest configuration file, resulting in one configuration file.

```
default: /etc/pgbackrest/conf.d
example: --config-include-path=/conf/pgbackrest/conf.d
```

**Config Path Option ( `-config-path` )** Base path of pgBackRest configuration files.

This setting is used to override the default base path setting for the `-config` and `-config-include-path` options unless they are explicitly set on the command-line.

For example, passing only `-config-path=/conf/pgbackrest` results in the `-config` default being set to `/conf/pgbackrest/pgbackrest.conf` and the `-config-include-path` default being set to `/conf/pgbackrest/conf.d`.

```
default: /etc/pgbackrest
example: --config-path=/conf/pgbackrest
```

**Database Timeout Option ( `-db-timeout` )** Database query timeout.

Sets the timeout, in seconds, for queries against the database. This includes the `pg_start_backup()` and `pg_stop_backup()` functions which can each take a substantial amount of time. Because of this the timeout should be kept high unless you know that these functions will return quickly (i.e. if you have set `startfast=y` and you know that the database cluster will not generate many WAL segments during the backup).

```
default: 1800
allowed: 0.1-604800
example: --db-timeout=600
```

**Lock Path Option ( `-lock-path` )** Path where lock files are stored.

The lock path provides a location for pgBackRest to create lock files to prevent conflicting operations from being run concurrently.

```
default: /tmp/pgbackrest
example: --lock-path=/backup/db/lock
```

**Neutral Umask Option ( `-neutral-umask` )** Use a neutral umask.

Sets the umask to 0000 so modes in the repository are created in a sensible way. The default directory mode is 0750 and default file mode is 0640. The lock and log directories set the directory and file mode to 0770 and 0660 respectively.

To use the executing user's umask instead specify `neutral-umask=n` in the config file or `-no-neutral-umask` on the command line.

```
default: y
example: --no-neutral-umask
```

**Protocol Timeout Option ( `-protocol-timeout` )** Protocol timeout.

Sets the timeout, in seconds, that the local or remote process will wait for a new message to be received on the protocol layer. This prevents processes from waiting indefinitely for a message. The `protocol-timeout` option must be greater than the `db-timeout` option.

```
default: 1830
allowed: 0.1-604800
example: --protocol-timeout=630
```

**Stanza Option ( `-stanza` )** Defines the stanza.

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as `app` or `dw`, rather than the local cluster name, such as `main` or `prod`.

```
example: --stanza=main
```

## Log Options

**Console Log Level Option ( `-log-level-console` )** Level for console logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-console=error
```

**File Log Level Option ( `-log-level-file` )** Level for file logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
```

```
example: --log-level-file=debug
```

**Std Error Log Level Option ( `-log-level-stderr` )** Level for stderr logging.

Specifies which log levels will output to stderr rather than stdout (specified by `log-level-console`). The timestamp and process will not be output to stderr.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-stderr=error
```

**Log Path Option ( `-log-path` )** Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
```

```
example: --log-path=/backup/db/log
```

**Log Subprocesses Option ( `-log-subprocess` )** Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file`.

```
default: n
example: --log-subprocess
```

**Log Timestamp Option ( `-log-timestamp` )** Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
example: --no-log-timestamp
```

## Repository Options

**Repository Cipher Type Option ( `-repo-cipher-type` )** Cipher used to encrypt the repository.

The following repository types are supported:

- none - The repository is not encrypted
- aes-256-cbc - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. S3) supports encryption.

```
default: none
example: --repo1-cipher-type=aes-256-cbc
```

**Repository Path Option ( `-repo-path` )** Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: --repo1-path=/backup/db/backrest
```

**S3 Repository Bucket Option ( `-repo-s3-bucket` )** S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting `repo-path=` but it is usually best to specify a prefix, such as `/repo`, so logs and other AWS generated content can also be stored in the bucket.

```
example: --repo1-s3-bucket=pg-backup
```

**S3 SSL CA File Option ( `-repo-s3-ca-file` )** S3 SSL CA File.

Use a CA file other than the system default.

```
example: --repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

**S3 SSL CA Path Option ( `-repo-s3-ca-path` )** S3 SSL CA Path.

Use a CA path other than the system default.

```
example: --repo1-s3-ca-path=/etc/pki/tls/certs
```

**S3 Repository Endpoint Option ( `-repo-s3-endpoint` )** S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: --repo1-s3-endpoint=s3.amazonaws.com
```

**S3 Repository Host Option ( `-repo-s3-host` )** S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: --repo1-s3-host=127.0.0.1
```

**S3 Repository Region Option ( `-repo-s3-region` )** S3 repository region.

The AWS region where the bucket was created.

```
example: --repo1-s3-region=us-east-1
```

**S3 Repository Verify SSL Option ( `-repo-s3-verify-ssl` )** Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
```

```
example: --no-repo1-s3-verify-ssl
```

**Repository Type Option ( `-repo-type` )** Type of storage used for the repository.

The following repository types are supported:

- cifs - Like posix , but disables links and directory fsyncs
- posix - Posix-compliant file systems
- s3 - AWS Simple Storage Service

```
default: posix
```

```
example: --repo1-type=cifs
```

## Stanza Options

**PostgreSQL Host Option ( `-pg-host` )** PostgreSQL host for operating remotely via SSH.

Used for backups where the PostgreSQL host is different from the repository host.

```
example: --pg1-host=db.domain.com
```

Deprecated Name: db-host

**PostgreSQL Host Command Option ( `-pg-host-cmd` )** pgBackRest exe path on the PostgreSQL host.

Required only if the path to pgbackrest is different on the local and PostgreSQL hosts. If not defined, the database host exe path will be set the same as the local exe path.

```
example: --pg1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: db-cmd

**PostgreSQL Host Configuration Option ( `-pg-host-config` )** pgBackRest database host configuration file.

Sets the location of the configuration file on the PostgreSQL host. This is only required if the PostgreSQL host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
```

```
example: --pg1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: db-config



**PostgreSQL Host Configuration Include Path Option ( `-pg-host-config-include-path` )** pgBackRest database host configuration include path.

Sets the location of the configuration include path on the PostgreSQL host. This is only required if the PostgreSQL host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: --pg1-host-config-include-path=/conf/pgbackrest/conf.d
```

**PostgreSQL Host Configuration Path Option ( `-pg-host-config-path` )** pgBackRest database host configuration path.

Sets the location of the configuration path on the PostgreSQL host. This is only required if the PostgreSQL host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --pg1-host-config-path=/conf/pgbackrest
```

**PostgreSQL Host Port Option ( `-pg-host-port` )** PostgreSQL host port when pg-host is set.

Use this option to specify a non-default port for the PostgreSQL host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --pg1-host-port=25
```

Deprecated Name: db-ssh-port

**PostgreSQL Host User Option ( `-pg-host-user` )** PostgreSQL host logon user when pg-host is set.

This user will also own the remote pgBackRest process and will initiate connections to PostgreSQL . For this to work correctly the user should be the PostgreSQL database cluster owner which is generally postgres , the default.

```
default: postgres
example: --pg1-host-user=db_owner
```

Deprecated Name: db-user

**PostgreSQL Path Option ( `-pg-path` )** PostgreSQL data directory.

This should be the same as the data\_directory setting in postgresql.conf . Even though this value can be read from postgresql.conf or PostgreSQL it is prudent to set it in case those resources are not available during a restore or offline backup scenario.

The pg-path option is tested against the value reported by PostgreSQL on every online backup so it should always be current.

```
example: --pg1-path=/data/db
```

Deprecated Name: db-path

**PostgreSQL Port Option ( `-pg-port` )** PostgreSQL port.

Port that PostgreSQL is running on. This usually does not need to be specified as most PostgreSQL clusters run on the default port.

```
default: 5432
allowed: 0-65535
example: --pg1-port=6543
```

Deprecated Name: db-port

**PostgreSQL Socket Path Option ( `-pg-socket-path` )** PostgreSQL unix socket path.

The unix socket directory that was specified when PostgreSQL was started. pgBackRest will automatically look in the standard location for your OS so there is usually no need to specify this setting unless the socket directory was explicitly modified with the unix\_socket\_directory setting in postgresql.conf .

```
allowed: 0-65535
example: --pg1-socket-path=/var/run/postgresql
```

Deprecated Name: db-socket-path

## Stanza Delete Command ( stanza-delete )

The stanza-delete command removes data in the repository associated with a stanza.

### WARNING:

Use this command with caution — it will permanently remove all backups and archives from the pgBackRest repository for the specified stanza.

To delete a stanza:

- Shut down the PostgreSQL cluster associated with the stanza (or use `-force` to override).
- Run the stop command on the repository host.
- Run the stanza-delete command on the repository host.

Once the command successfully completes, it is the responsibility of the user to remove the stanza from all pgBackRest configuration files.

## Command Options

**Force Option ( `-force` )** Force stanza delete.

If PostgreSQL is still running for the stanza, then this option can be used to force the stanza to be deleted from the repository.

```
default: n
example: --no-force
```

## General Options

**Buffer Size Option ( `-buffer-size` )** Buffer size for file operations.

Set the buffer size used for copy, compress, and uncompress functions. A maximum of 3 buffers will be in use at a time per process. An additional maximum of 256K per process may be used for zlib buffers.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024. For example, the case-insensitive value 32k (or 32KB) can be used instead of 32768.

Allowed values, in bytes, are 16384 , 32768 , 65536 , 131072 , 262144 , 524288 , 1048576 , 2097152 , 4194304 , 8388608 , and 16777216 .

```
default: 4194304
example: --buffer-size=32K
```

**SSH client command Option ( `-cmd-ssh` )** Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the ssh executable is not in \$PATH.

```
default: ssh
example: --cmd-ssh=/usr/bin/ssh
```

**Compress Level Option ( `-compress-level` )** Compression level for stored files.

Sets the zlib level to be used for file compression when `compress=y` .

```
default: 6
allowed: 0-9
example: --compress-level=9
```

**Network Compress Level Option ( `-compress-level-network` )** Compression level for network transfer when `compress=n` .

Sets the zlib level to be used for protocol compression when `compress=n` and the database cluster is not on the same host as the repository. Protocol compression is used to reduce network traffic but can be disabled by setting `compress-level-network=0` . When `compress=y` the `compress-level-network` setting is ignored and `compress-level` is used instead so that the file is only compressed once. SSH compression is always disabled.

```
default: 3
allowed: 0-9
example: --compress-level-network=1
```

**Config Option ( `-config` )** pgBackRest configuration file.

Use this option to specify a different configuration file than the default.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --config=/conf/pgbackrest/pgbackrest.conf
```

**Config Include Path Option ( `-config-include-path` )** Path to additional pgBackRest configuration files.

Configuration files existing in the specified location with extension `.conf` will be concatenated with the pgBackRest configuration file, resulting in one configuration file.

```
default: /etc/pgbackrest/conf.d
example: --config-include-path=/conf/pgbackrest/conf.d
```

**Config Path Option ( `-config-path` )** Base path of pgBackRest configuration files.

This setting is used to override the default base path setting for the `-config` and `-config-include-path` options unless they are explicitly set on the command-line.

For example, passing only `-config-path=/conf/pgbackrest` results in the `-config` default being set to `/conf/pgbackrest/pgbackrest.conf` and the `-config-include-path` default being set to `/conf/pgbackrest/conf.d`.

```
default: /etc/pgbackrest
example: --config-path=/conf/pgbackrest
```

**Database Timeout Option ( `-db-timeout` )** Database query timeout.

Sets the timeout, in seconds, for queries against the database. This includes the `pg_start_backup()` and `pg_stop_backup()` functions which can each take a substantial amount of time. Because of this the timeout should be kept high unless you know that these functions will return quickly (i.e. if you have set `startfast=y` and you know that the database cluster will not generate many WAL segments during the backup).

```
default: 1800
allowed: 0.1-604800
example: --db-timeout=600
```

**Lock Path Option ( `-lock-path` )** Path where lock files are stored.

The lock path provides a location for pgBackRest to create lock files to prevent conflicting operations from being run concurrently.

```
default: /tmp/pgbackrest
example: --lock-path=/backup/db/lock
```

**Neutral Umask Option ( `-neutral-umask` )** Use a neutral umask.

Sets the umask to 0000 so modes in the repository are created in a sensible way. The default directory mode is 0750 and default file mode is 0640. The lock and log directories set the directory and file mode to 0770 and 0660 respectively.

To use the executing user's umask instead specify `neutral-umask=n` in the config file or `-no-neutral-umask` on the command line.

```
default: y
example: --no-neutral-umask
```

**Protocol Timeout Option ( `-protocol-timeout` )** Protocol timeout.

Sets the timeout, in seconds, that the local or remote process will wait for a new message to be received on the protocol layer. This prevents processes from waiting indefinitely for a message. The `protocol-timeout` option must be greater than the `db-timeout` option.

```
default: 1830
allowed: 0.1-604800
example: --protocol-timeout=630
```

**Stanza Option ( `-stanza` )** Defines the stanza.

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as app or dw, rather than the local cluster name, such as main or prod.

```
example: --stanza=main
```

## Log Options

**Console Log Level Option ( `-log-level-console` )** Level for console logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-console=error
```

**File Log Level Option ( `-log-level-file` )** Level for file logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
```

```
example: --log-level-file=debug
```

**Std Error Log Level Option ( `-log-level-stderr` )** Level for stderr logging.

Specifies which log levels will output to stderr rather than stdout (specified by `log-level-console`). The timestamp and process will not be output to stderr.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-stderr=error
```

**Log Path Option ( `-log-path` )** Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
example: --log-path=/backup/db/log
```

**Log Subprocesses Option ( `-log-subprocess` )** Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file` .

```
default: n
example: --log-subprocess
```

**Log Timestamp Option ( `-log-timestamp` )** Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
example: --no-log-timestamp
```

## Repository Options

**Repository Cipher Type Option ( `-repo-cipher-type` )** Cipher used to encrypt the repository.

The following repository types are supported:

- none - The repository is not encrypted
- aes-256-cbc - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. S3) supports encryption.

```
default: none
example: --repo1-cipher-type=aes-256-cbc
```

**Repository Path Option ( `-repo-path` )** Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: --repo1-path=/backup/db/backrest
```

**S3 Repository Bucket Option ( `-repo-s3-bucket` )** S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting `repo-path=` but it is usually best to specify a prefix, such as `/repo` , so logs and other AWS generated content can also be stored in the bucket.

```
example: --repo1-s3-bucket=pg-backup
```

**S3 SSL CA File Option ( `-repo-s3-ca-file` )** S3 SSL CA File.

Use a CA file other than the system default.

```
example: --repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

### **S3 SSL CA Path Option ( `--repo-s3-ca-path` )** S3 SSL CA Path.

Use a CA path other than the system default.

```
example: --repo1-s3-ca-path=/etc/pki/tls/certs
```

### **S3 Repository Endpoint Option ( `--repo-s3-endpoint` )** S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: --repo1-s3-endpoint=s3.amazonaws.com
```

### **S3 Repository Host Option ( `--repo-s3-host` )** S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: --repo1-s3-host=127.0.0.1
```

### **S3 Repository Region Option ( `--repo-s3-region` )** S3 repository region.

The AWS region where the bucket was created.

```
example: --repo1-s3-region=us-east-1
```

### **S3 Repository Verify SSL Option ( `--repo-s3-verify-ssl` )** Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
```

```
example: --no-repo1-s3-verify-ssl
```

### **Repository Type Option ( `--repo-type` )** Type of storage used for the repository.

The following repository types are supported:

- cifs - Like posix , but disables links and directory fsyncs
- posix - Posix-compliant file systems
- s3 - AWS Simple Storage Service

```
default: posix
```

```
example: --repo1-type=cifs
```

## **Stanza Options**

### **PostgreSQL Host Option ( `--pg-host` )** PostgreSQL host for operating remotely via SSH.

Used for backups where the PostgreSQL host is different from the repository host.

```
example: --pg1-host=db.domain.com
```

Deprecated Name: db-host

### **PostgreSQL Host Command Option ( `--pg-host-cmd` )** pgBackRest exe path on the PostgreSQL host.

Required only if the path to pgbackrest is different on the local and PostgreSQL hosts. If not defined, the database host exe path will be set the same as the local exe path.

```
example: --pg1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: db-cmd

**PostgreSQL Host Configuration Option ( `-pg-host-config` )** pgBackRest database host configuration file.

Sets the location of the configuration file on the PostgreSQL host. This is only required if the PostgreSQL host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --pg1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: db-config

**PostgreSQL Host Configuration Include Path Option ( `-pg-host-config-include-path` )** pgBackRest database host configuration include path.

Sets the location of the configuration include path on the PostgreSQL host. This is only required if the PostgreSQL host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: --pg1-host-config-include-path=/conf/pgbackrest/conf.d
```

**PostgreSQL Host Configuration Path Option ( `-pg-host-config-path` )** pgBackRest database host configuration path.

Sets the location of the configuration path on the PostgreSQL host. This is only required if the PostgreSQL host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --pg1-host-config-path=/conf/pgbackrest
```

**PostgreSQL Host Port Option ( `-pg-host-port` )** PostgreSQL host port when pg-host is set.

Use this option to specify a non-default port for the PostgreSQL host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --pg1-host-port=25
```

Deprecated Name: db-ssh-port

**PostgreSQL Host User Option ( `-pg-host-user` )** PostgreSQL host logon user when pg-host is set.

This user will also own the remote pgBackRest process and will initiate connections to PostgreSQL . For this to work correctly the user should be the PostgreSQL database cluster owner which is generally postgres , the default.

```
default: postgres
example: --pg1-host-user=db_owner
```

Deprecated Name: db-user

**PostgreSQL Path Option ( `-pg-path` )** PostgreSQL data directory.

This should be the same as the `data_directory` setting in `postgresql.conf` . Even though this value can be read from `postgresql.conf` or PostgreSQL it is prudent to set it in case those resources are not available during a restore or offline backup scenario.

The `pg-path` option is tested against the value reported by PostgreSQL on every online backup so it should always be current.

```
example: --pg1-path=/data/db
```

Deprecated Name: db-path

**PostgreSQL Port Option ( `-pg-port` )** PostgreSQL port.

Port that PostgreSQL is running on. This usually does not need to be specified as most PostgreSQL clusters run on the default port.

```
default: 5432
allowed: 0-65535
example: --pg1-port=6543
```

Deprecated Name: db-port

**PostgreSQL Socket Path Option ( `-pg-socket-path` )** PostgreSQL unix socket path.

The unix socket directory that was specified when PostgreSQL was started. pgBackRest will automatically look in the standard location for your OS so there is usually no need to specify this setting unless the socket directory was explicitly modified with the `unix_socket_directory` setting in `postgresql.conf` .

```
allowed: 0-65535
example: --pg1-socket-path=/var/run/postgresql
```

Deprecated Name: `db-socket-path`

## Stanza Upgrade Command ( `stanza-upgrade` )

Immediately after upgrading PostgreSQL to a newer major version, the `pg-path` for all pgBackRest configurations must be set to the new database location and the `stanza-upgrade` run on the repository host. If the database is offline use the `-no-online` option.

### Command Options

**Backup from Standby Option ( `-backup-standby` )** Backup from the standby cluster.

Enable backup from standby to reduce load on the primary cluster. This option requires that both the primary and standby hosts be configured.

```
default: n
example: --backup-standby
```

**Online Option ( `-online` )** Update an online cluster.

Specifying `-no-online` prevents pgBackRest from connecting to PostgreSQL when upgrading the stanza.

```
default: y
example: --no-online
```

### General Options

**Buffer Size Option ( `-buffer-size` )** Buffer size for file operations.

Set the buffer size used for copy, compress, and uncompress functions. A maximum of 3 buffers will be in use at a time per process. An additional maximum of 256K per process may be used for zlib buffers.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024. For example, the case-insensitive value `32k` (or `32KB`) can be used instead of `32768`.

Allowed values, in bytes, are 16384 , 32768 , 65536 , 131072 , 262144 , 524288 , 1048576 , 2097152 , 4194304 , 8388608 , and 16777216 .

```
default: 4194304
example: --buffer-size=32K
```

**SSH client command Option ( `-cmd-ssh` )** Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the `ssh` executable is not in `$PATH`.

```
default: ssh
example: --cmd-ssh=/usr/bin/ssh
```

**Compress Level Option ( `-compress-level` )** Compression level for stored files.

Sets the zlib level to be used for file compression when `compress=y` .

```
default: 6
allowed: 0-9
example: --compress-level=9
```



**Network Compress Level Option ( `-compress-level-network` )** Compression level for network transfer when `compress=n` .

Sets the zlib level to be used for protocol compression when `compress=n` and the database cluster is not on the same host as the repository. Protocol compression is used to reduce network traffic but can be disabled by setting `compress-level-network=0` . When `compress=y` the `compress-level-network` setting is ignored and `compress-level` is used instead so that the file is only compressed once. SSH compression is always disabled.

```
default: 3
allowed: 0-9
example: --compress-level-network=1
```

**Config Option ( `-config` )** pgBackRest configuration file.

Use this option to specify a different configuration file than the default.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --config=/conf/pgbackrest/pgbackrest.conf
```

**Config Include Path Option ( `-config-include-path` )** Path to additional pgBackRest configuration files.

Configuration files existing in the specified location with extension `.conf` will be concatenated with the pgBackRest configuration file, resulting in one configuration file.

```
default: /etc/pgbackrest/conf.d
example: --config-include-path=/conf/pgbackrest/conf.d
```

**Config Path Option ( `-config-path` )** Base path of pgBackRest configuration files.

This setting is used to override the default base path setting for the `-config` and `-config-include-path` options unless they are explicitly set on the command-line.

For example, passing only `-config-path=/conf/pgbackrest` results in the `-config` default being set to `/conf/pgbackrest/pgbackrest.conf` and the `-config-include-path` default being set to `/conf/pgbackrest/conf.d` .

```
default: /etc/pgbackrest
example: --config-path=/conf/pgbackrest
```

**Database Timeout Option ( `-db-timeout` )** Database query timeout.

Sets the timeout, in seconds, for queries against the database. This includes the `pg_start_backup()` and `pg_stop_backup()` functions which can each take a substantial amount of time. Because of this the timeout should be kept high unless you know that these functions will return quickly (i.e. if you have set `startfast=y` and you know that the database cluster will not generate many WAL segments during the backup).

```
default: 1800
allowed: 0.1-604800
example: --db-timeout=600
```

**Lock Path Option ( `-lock-path` )** Path where lock files are stored.

The lock path provides a location for pgBackRest to create lock files to prevent conflicting operations from being run concurrently.

```
default: /tmp/pgbackrest
example: --lock-path=/backup/db/lock
```

**Neutral Umask Option ( `-neutral-umask` )** Use a neutral umask.

Sets the umask to 0000 so modes in the repository are created in a sensible way. The default directory mode is 0750 and default file mode is 0640. The lock and log directories set the directory and file mode to 0770 and 0660 respectively.

To use the executing user's umask instead specify `neutral-umask=n` in the config file or `-no-neutral-umask` on the command line.

```
default: y
example: --no-neutral-umask
```

**Protocol Timeout Option ( `-protocol-timeout` )** Protocol timeout.

Sets the timeout, in seconds, that the local or remote process will wait for a new message to be received on the protocol layer. This prevents processes from waiting indefinitely for a message. The `protocol-timeout` option must be greater than the `db-timeout` option.

```
default: 1830
allowed: 0.1-604800
example: --protocol-timeout=630
```

**Stanza Option ( `-stanza` )** Defines the stanza.

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as `app` or `dw`, rather than the local cluster name, such as `main` or `prod`.

```
example: --stanza=main
```

## Log Options

**Console Log Level Option ( `-log-level-console` )** Level for console logging.

The following log levels are supported:

- `off` - No logging at all (not recommended)
- `error` - Log only errors
- `warn` - Log warnings and errors
- `info` - Log info, warnings, and errors
- `detail` - Log detail, info, warnings, and errors
- `debug` - Log debug, detail, info, warnings, and errors
- `trace` - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
example: --log-level-console=error
```

**File Log Level Option ( `-log-level-file` )** Level for file logging.

The following log levels are supported:

- `off` - No logging at all (not recommended)
- `error` - Log only errors
- `warn` - Log warnings and errors
- `info` - Log info, warnings, and errors
- `detail` - Log detail, info, warnings, and errors
- `debug` - Log debug, detail, info, warnings, and errors
- `trace` - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
example: --log-level-file=debug
```

**Std Error Log Level Option ( `-log-level-stderr` )** Level for stderr logging.

Specifies which log levels will output to `stderr` rather than `stdout` (specified by `log-level-console`). The timestamp and process will not be output to `stderr`.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
example: --log-level-stderr=error
```

**Log Path Option ( `-log-path` )** Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
example: --log-path=/backup/db/log
```

**Log Subprocesses Option ( `-log-subprocess` )** Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file` .

```
default: n
example: --log-subprocess
```

**Log Timestamp Option ( `-log-timestamp` )** Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
example: --no-log-timestamp
```

## Repository Options

**Repository Cipher Type Option ( `-repo-cipher-type` )** Cipher used to encrypt the repository.

The following repository types are supported:

- none - The repository is not encrypted
- aes-256-cbc - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. S3) supports encryption.

```
default: none
example: --repo1-cipher-type=aes-256-cbc
```

**Repository Path Option ( `-repo-path` )** Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: --repo1-path=/backup/db/backrest
```

**S3 Repository Bucket Option ( `-repo-s3-bucket` )** S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting `repo-path=` but it is usually best to specify a prefix, such as `/repo`, so logs and other AWS generated content can also be stored in the bucket.

```
example: --repo1-s3-bucket=pg-backup
```

**S3 SSL CA File Option ( `-repo-s3-ca-file` )** S3 SSL CA File.

Use a CA file other than the system default.

```
example: --repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

**S3 SSL CA Path Option ( `-repo-s3-ca-path` )** S3 SSL CA Path.

Use a CA path other than the system default.

```
example: --repo1-s3-ca-path=/etc/pki/tls/certs
```

**S3 Repository Endpoint Option ( `-repo-s3-endpoint` )** S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: --repo1-s3-endpoint=s3.amazonaws.com
```

**S3 Repository Host Option ( `-repo-s3-host` )** S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: --repo1-s3-host=127.0.0.1
```

**S3 Repository Region Option ( `-repo-s3-region` )** S3 repository region.

The AWS region where the bucket was created.

```
example: --repo1-s3-region=us-east-1
```

**S3 Repository Verify SSL Option ( `-repo-s3-verify-ssl` )** Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
```

```
example: --no-repo1-s3-verify-ssl
```

**Repository Type Option ( `-repo-type` )** Type of storage used for the repository.

The following repository types are supported:

- `cifs` - Like `posix`, but disables links and directory fsyncs
- `posix` - Posix-compliant file systems
- `s3` - AWS Simple Storage Service

```
default: posix
```

```
example: --repo1-type=cifs
```

## Stanza Options

**PostgreSQL Host Option ( `-pg-host` )** PostgreSQL host for operating remotely via SSH.

Used for backups where the PostgreSQL host is different from the repository host.

```
example: --pg1-host=db.domain.com
```

Deprecated Name: db-host

**PostgreSQL Host Command Option ( `-pg-host-cmd` )** pgBackRest exe path on the PostgreSQL host.

Required only if the path to pgbackrest is different on the local and PostgreSQL hosts. If not defined, the database host exe path will be set the same as the local exe path.

```
example: --pg1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: db-cmd

**PostgreSQL Host Configuration Option ( `-pg-host-config` )** pgBackRest database host configuration file.

Sets the location of the configuration file on the PostgreSQL host. This is only required if the PostgreSQL host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
```

```
example: --pg1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: db-config

**PostgreSQL Host Configuration Include Path Option ( `-pg-host-config-include-path` )** pgBackRest database host configuration include path.

Sets the location of the configuration include path on the PostgreSQL host. This is only required if the PostgreSQL host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
```

```
example: --pg1-host-config-include-path=/conf/pgbackrest/conf.d
```

**PostgreSQL Host Configuration Path Option ( `-pg-host-config-path` )** pgBackRest database host configuration path.

Sets the location of the configuration path on the PostgreSQL host. This is only required if the PostgreSQL host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
```

```
example: --pg1-host-config-path=/conf/pgbackrest
```

**PostgreSQL Host Port Option ( `-pg-host-port` )** PostgreSQL host port when pg-host is set.

Use this option to specify a non-default port for the PostgreSQL host protocol. Currently only SSH is supported

```
allowed: 0-65535
```

```
example: --pg1-host-port=25
```

Deprecated Name: db-ssh-port

**PostgreSQL Host User Option ( `-pg-host-user` )** PostgreSQL host logon user when pg-host is set.

This user will also own the remote pgBackRest process and will initiate connections to PostgreSQL . For this to work correctly the user should be the PostgreSQL database cluster owner which is generally postgres , the default.

```
default: postgres
```

```
example: --pg1-host-user=db_owner
```

Deprecated Name: db-user

**PostgreSQL Path Option ( `-pg-path` )** PostgreSQL data directory.

This should be the same as the `data_directory` setting in `postgresql.conf` . Even though this value can be read from `postgresql.conf` or PostgreSQL it is prudent to set it in case those resources are not available during a restore or offline backup scenario.

The `pg-path` option is tested against the value reported by PostgreSQL on every online backup so it should always be current.

```
example: --pg1-path=/data/db
```

Deprecated Name: `db-path`

**PostgreSQL Port Option ( `-pg-port` )** PostgreSQL port.

Port that PostgreSQL is running on. This usually does not need to be specified as most PostgreSQL clusters run on the default port.

```
default: 5432
allowed: 0-65535
example: --pg1-port=6543
```

Deprecated Name: `db-port`

**PostgreSQL Socket Path Option ( `-pg-socket-path` )** PostgreSQL unix socket path.

The unix socket directory that was specified when PostgreSQL was started. `pgBackRest` will automatically look in the standard location for your OS so there is usually no need to specify this setting unless the socket directory was explicitly modified with the `unix_socket_directory` setting in `postgresql.conf` .

```
allowed: 0-65535
example: --pg1-socket-path=/var/run/postgresql
```

Deprecated Name: `db-socket-path`

## Start Command ( `start` )

If the `pgBackRest` processes were previously stopped using the `stop` command then they can be started again using the `start` command. Note that this will not immediately start up any `pgBackRest` processes but they are allowed to run.

## General Options

**SSH client command Option ( `-cmd-ssh` )** Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the `ssh` executable is not in `$PATH`.

```
default: ssh
example: --cmd-ssh=/usr/bin/ssh
```

**Config Option ( `-config` )** `pgBackRest` configuration file.

Use this option to specify a different configuration file than the default.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --config=/conf/pgbackrest/pgbackrest.conf
```

**Config Include Path Option ( `-config-include-path` )** Path to additional `pgBackRest` configuration files.

Configuration files existing in the specified location with extension `.conf` will be concatenated with the `pgBackRest` configuration file, resulting in one configuration file.

```
default: /etc/pgbackrest/conf.d
example: --config-include-path=/conf/pgbackrest/conf.d
```

**Config Path Option ( `-config-path` )** Base path of pgBackRest configuration files.

This setting is used to override the default base path setting for the `-config` and `-config-include-path` options unless they are explicitly set on the command-line.

For example, passing only `-config-path=/conf/pgbackrest` results in the `-config` default being set to `/conf/pgbackrest/pgbackrest.conf` and the `-config-include-path` default being set to `/conf/pgbackrest/conf.d`.

```
default: /etc/pgbackrest
example: --config-path=/conf/pgbackrest
```

**Lock Path Option ( `-lock-path` )** Path where lock files are stored.

The lock path provides a location for pgBackRest to create lock files to prevent conflicting operations from being run concurrently.

```
default: /tmp/pgbackrest
example: --lock-path=/backup/db/lock
```

**Neutral Umask Option ( `-neutral-umask` )** Use a neutral umask.

Sets the umask to 0000 so modes in the repository are created in a sensible way. The default directory mode is 0750 and default file mode is 0640. The lock and log directories set the directory and file mode to 0770 and 0660 respectively.

To use the executing user's umask instead specify `neutral-umask=n` in the config file or `-no-neutral-umask` on the command line.

```
default: y
example: --no-neutral-umask
```

**Stanza Option ( `-stanza` )** Defines the stanza.

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as `app` or `dw`, rather than the local cluster name, such as `main` or `prod`.

```
example: --stanza=main
```

## Log Options

**Console Log Level Option ( `-log-level-console` )** Level for console logging.

The following log levels are supported:

- `off` - No logging at all (not recommended)
- `error` - Log only errors
- `warn` - Log warnings and errors
- `info` - Log info, warnings, and errors
- `detail` - Log detail, info, warnings, and errors
- `debug` - Log debug, detail, info, warnings, and errors
- `trace` - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
example: --log-level-console=error
```

**File Log Level Option ( `-log-level-file` )** Level for file logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
```

```
example: --log-level-file=debug
```

**Std Error Log Level Option ( `-log-level-stderr` )** Level for stderr logging.

Specifies which log levels will output to stderr rather than stdout (specified by `log-level-console` ). The timestamp and process will not be output to stderr .

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-stderr=error
```

**Log Path Option ( `-log-path` )** Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
```

```
example: --log-path=/backup/db/log
```

**Log Subprocesses Option ( `-log-subprocess` )** Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file` .

```
default: n
```

```
example: --log-subprocess
```

**Log Timestamp Option ( `-log-timestamp` )** Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
```

```
example: --no-log-timestamp
```



## Repository Options

**Repository Cipher Type Option ( `--repo-cipher-type` )** Cipher used to encrypt the repository.

The following repository types are supported:

- none - The repository is not encrypted
- aes-256-cbc - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. S3) supports encryption.

```
default: none
example: --repo1-cipher-type=aes-256-cbc
```

**Repository Host Option ( `--repo-host` )** Repository host when operating remotely via SSH.

Make sure that trusted SSH authentication is configured between the PostgreSQL host and the repository host.

When backing up and archiving to a locally mounted filesystem this setting is not required.

```
example: --repo1-host=repo1.domain.com
```

Deprecated Name: backup-host

**Repository Host Command Option ( `--repo-host-cmd` )** pgBackRest exe path on the repository host.

Required only if the path to pgbackrest is different on the local and repository hosts. If not defined, the repository host exe path will be set the same as the local exe path.

```
example: --repo1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: backup-cmd

**Repository Host Configuration Option ( `--repo-host-config` )** pgBackRest repository host configuration file.

Sets the location of the configuration file on the repository host. This is only required if the repository host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --repo1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: backup-config

**Repository Host Configuration Include Path Option ( `--repo-host-config-include-path` )** pgBackRest repository host configuration include path.

Sets the location of the configuration include path on the repository host. This is only required if the repository host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: --repo1-host-config-include-path=/conf/pgbackrest/conf.d
```

**Repository Host Configuration Path Option ( `--repo-host-config-path` )** pgBackRest repository host configuration path.

Sets the location of the configuration path on the repository host. This is only required if the repository host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --repo1-host-config-path=/conf/pgbackrest
```

**Repository Host Port Option ( `--repo-host-port` )** Repository host port when repo-host is set.

Use this option to specify a non-default port for the repository host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --repo1-host-port=25
```

Deprecated Name: backup-ssh-port

**Repository Host User Option ( `-repo-host-user` )** Repository host user when repo-host is set.

Defines the user that will be used for operations on the repository host. Preferably this is not the postgres user but rather some other user like pgbackrest . If PostgreSQL runs on the repository host the postgres user can be placed in the pgbackrest group so it has read permissions on the repository without being able to damage the contents accidentally.

```
default: pgbackrest
example: --repo1-host-user=repo-user
```

Deprecated Name: backup-user

**Repository Path Option ( `-repo-path` )** Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: --repo1-path=/backup/db/backrest
```

**S3 Repository Bucket Option ( `-repo-s3-bucket` )** S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting repo-path=/ but it is usually best to specify a prefix, such as /repo , so logs and other AWS generated content can also be stored in the bucket.

```
example: --repo1-s3-bucket=pg-backup
```

**S3 SSL CA File Option ( `-repo-s3-ca-file` )** S3 SSL CA File.

Use a CA file other than the system default.

```
example: --repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

**S3 SSL CA Path Option ( `-repo-s3-ca-path` )** S3 SSL CA Path.

Use a CA path other than the system default.

```
example: --repo1-s3-ca-path=/etc/pki/tls/certs
```

**S3 Repository Endpoint Option ( `-repo-s3-endpoint` )** S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: --repo1-s3-endpoint=s3.amazonaws.com
```

**S3 Repository Host Option ( `-repo-s3-host` )** S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: --repo1-s3-host=127.0.0.1
```

**S3 Repository Region Option ( `-repo-s3-region` )** S3 repository region.

The AWS region where the bucket was created.

```
example: --repo1-s3-region=us-east-1
```

**S3 Repository Verify SSL Option ( `--repo-s3-verify-ssl` )** Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
example: --no-repo1-s3-verify-ssl
```

**Repository Type Option ( `--repo-type` )** Type of storage used for the repository.

The following repository types are supported:

- cifs - Like posix , but disables links and directory fsyncs
- posix - Posix-compliant file systems
- s3 - AWS Simple Storage Service

```
default: posix
example: --repo1-type=cifs
```

### Stanza Options

**PostgreSQL Host Option ( `--pg-host` )** PostgreSQL host for operating remotely via SSH.

Used for backups where the PostgreSQL host is different from the repository host.

```
example: --pg1-host=db.domain.com
```

Deprecated Name: db-host

**PostgreSQL Host Command Option ( `--pg-host-cmd` )** pgBackRest exe path on the PostgreSQL host.

Required only if the path to pgbackrest is different on the local and PostgreSQL hosts. If not defined, the database host exe path will be set the same as the local exe path.

```
example: --pg1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: db-cmd

**PostgreSQL Host Configuration Option ( `--pg-host-config` )** pgBackRest database host configuration file.

Sets the location of the configuration file on the PostgreSQL host. This is only required if the PostgreSQL host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --pg1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: db-config

**PostgreSQL Host Configuration Include Path Option ( `--pg-host-config-include-path` )** pgBackRest database host configuration include path.

Sets the location of the configuration include path on the PostgreSQL host. This is only required if the PostgreSQL host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: --pg1-host-config-include-path=/conf/pgbackrest/conf.d
```

**PostgreSQL Host Configuration Path Option ( `--pg-host-config-path` )** pgBackRest database host configuration path.

Sets the location of the configuration path on the PostgreSQL host. This is only required if the PostgreSQL host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --pg1-host-config-path=/conf/pgbackrest
```

**PostgreSQL Host Port Option ( `-pg-host-port` )** PostgreSQL host port when `pg-host` is set.

Use this option to specify a non-default port for the PostgreSQL host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --pg1-host-port=25
```

Deprecated Name: `db-ssh-port`

**PostgreSQL Host User Option ( `-pg-host-user` )** PostgreSQL host logon user when `pg-host` is set.

This user will also own the remote `pgBackRest` process and will initiate connections to PostgreSQL . For this to work correctly the user should be the PostgreSQL database cluster owner which is generally `postgres` , the default.

```
default: postgres
example: --pg1-host-user=db_owner
```

Deprecated Name: `db-user`

## Stop Command ( `stop` )

Does not allow any new `pgBackRest` processes to run. By default running processes will be allowed to complete successfully. Use the `-force` option to terminate running processes.

`pgBackRest` processes will return an error if they are run after the stop command completes.

## Command Options

**Force Option ( `-force` )** Force all `pgBackRest` processes to stop.

This option will send `TERM` signals to all running `pgBackRest` processes to effect a graceful but immediate shutdown. Note that this will also shutdown processes that were initiated on another system but have remotes running on the current system. For instance, if a backup was started on the backup server then running `stop -force` on the database server will shutdown the backup process on the backup server.

```
default: n
example: --force
```

## General Options

**SSH client command Option ( `-cmd-ssh` )** Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the `ssh` executable is not in `$PATH`.

```
default: ssh
example: --cmd-ssh=/usr/bin/ssh
```

**Config Option ( `-config` )** `pgBackRest` configuration file.

Use this option to specify a different configuration file than the default.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --config=/conf/pgbackrest/pgbackrest.conf
```

**Config Include Path Option ( `-config-include-path` )** Path to additional `pgBackRest` configuration files.

Configuration files existing in the specified location with extension `.conf` will be concatenated with the `pgBackRest` configuration file, resulting in one configuration file.

```
default: /etc/pgbackrest/conf.d
example: --config-include-path=/conf/pgbackrest/conf.d
```

**Config Path Option ( `-config-path` )** Base path of pgBackRest configuration files.

This setting is used to override the default base path setting for the `-config` and `-config-include-path` options unless they are explicitly set on the command-line.

For example, passing only `-config-path=/conf/pgbackrest` results in the `-config` default being set to `/conf/pgbackrest/pgbackrest.conf` and the `-config-include-path` default being set to `/conf/pgbackrest/conf.d`.

```
default: /etc/pgbackrest
example: --config-path=/conf/pgbackrest
```

**Lock Path Option ( `-lock-path` )** Path where lock files are stored.

The lock path provides a location for pgBackRest to create lock files to prevent conflicting operations from being run concurrently.

```
default: /tmp/pgbackrest
example: --lock-path=/backup/db/lock
```

**Neutral Umask Option ( `-neutral-umask` )** Use a neutral umask.

Sets the umask to 0000 so modes in the repository are created in a sensible way. The default directory mode is 0750 and default file mode is 0640. The lock and log directories set the directory and file mode to 0770 and 0660 respectively.

To use the executing user's umask instead specify `neutral-umask=n` in the config file or `-no-neutral-umask` on the command line.

```
default: y
example: --no-neutral-umask
```

**Stanza Option ( `-stanza` )** Defines the stanza.

A stanza is the configuration for a PostgreSQL database cluster that defines where it is located, how it will be backed up, archiving options, etc. Most db servers will only have one Postgres database cluster and therefore one stanza, whereas backup servers will have a stanza for every database cluster that needs to be backed up.

It is tempting to name the stanza after the primary cluster but a better name describes the databases contained in the cluster. Because the stanza name will be used for the primary and all replicas it is more appropriate to choose a name that describes the actual function of the cluster, such as `app` or `dw`, rather than the local cluster name, such as `main` or `prod`.

```
example: --stanza=main
```

## Log Options

**Console Log Level Option ( `-log-level-console` )** Level for console logging.

The following log levels are supported:

- `off` - No logging at all (not recommended)
- `error` - Log only errors
- `warn` - Log warnings and errors
- `info` - Log info, warnings, and errors
- `detail` - Log detail, info, warnings, and errors
- `debug` - Log debug, detail, info, warnings, and errors
- `trace` - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
example: --log-level-console=error
```

**File Log Level Option ( `-log-level-file` )** Level for file logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
```

```
example: --log-level-file=debug
```

**Std Error Log Level Option ( `-log-level-stderr` )** Level for stderr logging.

Specifies which log levels will output to stderr rather than stdout (specified by `log-level-console` ). The timestamp and process will not be output to stderr .

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: --log-level-stderr=error
```

**Log Path Option ( `-log-path` )** Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
```

```
example: --log-path=/backup/db/log
```

**Log Subprocesses Option ( `-log-subprocess` )** Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file` .

```
default: n
```

```
example: --log-subprocess
```

**Log Timestamp Option ( `-log-timestamp` )** Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
```

```
example: --no-log-timestamp
```

## Repository Options

**Repository Cipher Type Option ( `--repo-cipher-type` )** Cipher used to encrypt the repository.

The following repository types are supported:

- none - The repository is not encrypted
- aes-256-cbc - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. S3) supports encryption.

```
default: none
example: --repo1-cipher-type=aes-256-cbc
```

**Repository Host Option ( `--repo-host` )** Repository host when operating remotely via SSH.

Make sure that trusted SSH authentication is configured between the PostgreSQL host and the repository host.

When backing up and archiving to a locally mounted filesystem this setting is not required.

```
example: --repo1-host=repo1.domain.com
```

Deprecated Name: backup-host

**Repository Host Command Option ( `--repo-host-cmd` )** pgBackRest exe path on the repository host.

Required only if the path to pgbackrest is different on the local and repository hosts. If not defined, the repository host exe path will be set the same as the local exe path.

```
example: --repo1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: backup-cmd

**Repository Host Configuration Option ( `--repo-host-config` )** pgBackRest repository host configuration file.

Sets the location of the configuration file on the repository host. This is only required if the repository host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --repo1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: backup-config

**Repository Host Configuration Include Path Option ( `--repo-host-config-include-path` )** pgBackRest repository host configuration include path.

Sets the location of the configuration include path on the repository host. This is only required if the repository host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: --repo1-host-config-include-path=/conf/pgbackrest/conf.d
```

**Repository Host Configuration Path Option ( `--repo-host-config-path` )** pgBackRest repository host configuration path.

Sets the location of the configuration path on the repository host. This is only required if the repository host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --repo1-host-config-path=/conf/pgbackrest
```

**Repository Host Port Option ( `--repo-host-port` )** Repository host port when repo-host is set.

Use this option to specify a non-default port for the repository host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --repo1-host-port=25
```

Deprecated Name: backup-ssh-port

**Repository Host User Option ( `-repo-host-user` )** Repository host user when repo-host is set.

Defines the user that will be used for operations on the repository host. Preferably this is not the postgres user but rather some other user like pgbackrest . If PostgreSQL runs on the repository host the postgres user can be placed in the pgbackrest group so it has read permissions on the repository without being able to damage the contents accidentally.

```
default: pgbackrest
example: --repo1-host-user=repo-user
```

Deprecated Name: backup-user

**Repository Path Option ( `-repo-path` )** Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: --repo1-path=/backup/db/backrest
```

**S3 Repository Bucket Option ( `-repo-s3-bucket` )** S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting repo-path=/ but it is usually best to specify a prefix, such as /repo , so logs and other AWS generated content can also be stored in the bucket.

```
example: --repo1-s3-bucket=pg-backup
```

**S3 SSL CA File Option ( `-repo-s3-ca-file` )** S3 SSL CA File.

Use a CA file other than the system default.

```
example: --repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

**S3 SSL CA Path Option ( `-repo-s3-ca-path` )** S3 SSL CA Path.

Use a CA path other than the system default.

```
example: --repo1-s3-ca-path=/etc/pki/tls/certs
```

**S3 Repository Endpoint Option ( `-repo-s3-endpoint` )** S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: --repo1-s3-endpoint=s3.amazonaws.com
```

**S3 Repository Host Option ( `-repo-s3-host` )** S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: --repo1-s3-host=127.0.0.1
```

**S3 Repository Region Option ( `-repo-s3-region` )** S3 repository region.

The AWS region where the bucket was created.

```
example: --repo1-s3-region=us-east-1
```



**S3 Repository Verify SSL Option ( `--repo-s3-verify-ssl` )** Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
example: --no-repo1-s3-verify-ssl
```

**Repository Type Option ( `--repo-type` )** Type of storage used for the repository.

The following repository types are supported:

- cifs - Like posix , but disables links and directory fsyncs
- posix - Posix-compliant file systems
- s3 - AWS Simple Storage Service

```
default: posix
example: --repo1-type=cifs
```

### Stanza Options

**PostgreSQL Host Option ( `--pg-host` )** PostgreSQL host for operating remotely via SSH.

Used for backups where the PostgreSQL host is different from the repository host.

```
example: --pg1-host=db.domain.com
```

Deprecated Name: db-host

**PostgreSQL Host Command Option ( `--pg-host-cmd` )** pgBackRest exe path on the PostgreSQL host.

Required only if the path to pgbackrest is different on the local and PostgreSQL hosts. If not defined, the database host exe path will be set the same as the local exe path.

```
example: --pg1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: db-cmd

**PostgreSQL Host Configuration Option ( `--pg-host-config` )** pgBackRest database host configuration file.

Sets the location of the configuration file on the PostgreSQL host. This is only required if the PostgreSQL host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
example: --pg1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: db-config

**PostgreSQL Host Configuration Include Path Option ( `--pg-host-config-include-path` )** pgBackRest database host configuration include path.

Sets the location of the configuration include path on the PostgreSQL host. This is only required if the PostgreSQL host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: --pg1-host-config-include-path=/conf/pgbackrest/conf.d
```

**PostgreSQL Host Configuration Path Option ( `--pg-host-config-path` )** pgBackRest database host configuration path.

Sets the location of the configuration path on the PostgreSQL host. This is only required if the PostgreSQL host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: --pg1-host-config-path=/conf/pgbackrest
```

**PostgreSQL Host Port Option ( `-pg-host-port` )** PostgreSQL host port when pg-host is set.

Use this option to specify a non-default port for the PostgreSQL host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: --pg1-host-port=25
```

Deprecated Name: db-ssh-port

**PostgreSQL Host User Option ( `-pg-host-user` )** PostgreSQL host logon user when pg-host is set.

This user will also own the remote pgBackRest process and will initiate connections to PostgreSQL . For this to work correctly the user should be the PostgreSQL database cluster owner which is generally postgres , the default.

```
default: postgres
example: --pg1-host-user=db_owner
```

Deprecated Name: db-user

## Version Command ( `version` )

Displays installed pgBackRest version.

— title: Configuration Reference draft: false —

## Introduction

pgBackRest can be used entirely with command-line parameters but a configuration file is more practical for installations that are complex or set a lot of options. The default location for the configuration file is `/etc/pgbackrest/pgbackrest.conf` . If no file exists in that location then the old default of `/etc/pgbackrest.conf` will be checked.

## Archive Options ( `archive` )

The archive section defines options for the archive-push and archive-get commands.

### Asynchronous Archiving Option ( `-archive-async` )

Push/get WAL segments asynchronously.

Enables asynchronous operation for the archive-push and archive-get commands.

Asynchronous operation is more efficient because it can reuse connections and take advantage of parallelism. See the `spool-path` , `archive-get-queue-max` , and `archive-push-queue-max` options for more information.

```
default: n
example: archive-async=y
```

### Maximum Archive Get Queue Size Option ( `-archive-get-queue-max` )

Maximum size of the pgBackRest archive-get queue.

Specifies the maximum size of the archive-get queue when archive-async is enabled. The queue is stored in the `spool-path` and is used to speed providing WAL to PostgreSQL .

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024.

```
default: 134217728
allowed: 0-4503599627370496
example: archive-get-queue-max=1073741824
```

## Maximum Archive Push Queue Size Option ( `-archive-push-queue-max` )

Maximum size of the PostgreSQL archive queue.

After the limit is reached, the following will happen:

1. pgBackRest will notify PostgreSQL that the WAL was successfully archived, then **DROP IT** .
2. A warning will be output to the Postgres log.

If this occurs then the archive log stream will be interrupted and PITR will not be possible past that point. A new backup will be required to regain full restore capability.

In asynchronous mode the entire queue will be dropped to prevent spurts of WAL getting through before the queue limit is exceeded again.

The purpose of this feature is to prevent the log volume from filling up at which point Postgres will stop completely. Better to lose the backup than have PostgreSQL go down.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024.

```
allowed: 0-4503599627370496
```

```
example: archive-push-queue-max=1GB
```

Deprecated Name: `archive-queue-max`

## Archive Timeout Option ( `-archive-timeout` )

Archive timeout.

Set maximum time, in seconds, to wait for each WAL segment to reach the pgBackRest archive repository. The timeout applies to the check and backup commands when waiting for WAL segments required for backup consistency to be archived.

```
default: 60
```

```
allowed: 0.1-86400
```

```
example: archive-timeout=30
```

## Backup Options ( `backup` )

The backup section defines settings related to backup.

### Check Archive Option ( `-archive-check` )

Check that WAL segments are in the archive before backup completes.

Checks that all WAL segments required to make the backup consistent are present in the WAL archive. It's a good idea to leave this as the default unless you are using another method for archiving.

This option must be enabled if `archive-copy` is enabled.

```
default: y
```

```
example: archive-check=n
```

### Copy Archive Option ( `-archive-copy` )

Copy WAL segments needed for consistency to the backup.

This slightly paranoid option protects against corruption in the WAL segment archive by storing the WAL segments required for consistency directly in the backup. WAL segments are still stored in the archive so this option will use additional space.

On restore, the WAL segments will be present in `pg_xlog/pg_wal` and PostgreSQL will use them in preference to calling the `restore_command` .

The `archive-check` option must be enabled if `archive-copy` is enabled.

```
default: n
```

```
example: archive-copy=y
```

## Backup from Standby Option ( `-backup-standby` )

Backup from the standby cluster.

Enable backup from standby to reduce load on the primary cluster. This option requires that both the primary and standby hosts be configured.

```
default: n
example: backup-standby=y
```

## Page Checksums Option ( `-checksum-page` )

Validate data page checksums.

Directs pgBackRest to validate all data page checksums while backing up a cluster. This option is automatically enabled when data page checksums are enabled on the cluster.

Failures in checksum validation will not abort a backup. Rather, warnings will be emitted in the log (and to the console with default settings) and the list of invalid pages will be stored in the backup manifest.

```
example: checksum-page=n
```

## Path/File Exclusions Option ( `-exclude` )

Exclude paths/files from the backup.

All exclusions are relative to `$PGDATA` . If the exclusion ends with `/` then only files in the specified directory will be excluded, e.g. `-exclude=junk/` will exclude all files in the `$PGDATA/junk` directory but include the directory itself. If the exclusion does not end with `/` then the file may match the exclusion exactly or match with `/` appended to the exclusion, e.g. `-exclude=junk` will exclude the `$PGDATA/junk` directory and all the files it contains.

**Be careful using this feature – it is very easy to exclude something critical that will make the backup inconsistent. Be sure to test your restores!**

All excluded files will be logged at info level along with the exclusion rule. Be sure to audit the list of excluded files to ensure nothing unexpected is being excluded.

NOTE:

Exclusions are not honored on delta restores. Any files/directories that were excluded by the backup will be *removed* on delta restore.

This option should not be used to exclude PostgreSQL logs from a backup. Logs can be moved out of the `PGDATA` directory using the PostgreSQL `log_directory` setting, which has the benefit of allowing logs to be preserved after a restore.

Multiple exclusions may be specified on the command-line or in a configuration file.

```
example: exclude=junk/
```

## Manifest Save Threshold Option ( `-manifest-save-threshold` )

Manifest save threshold during backup.

Defines how often the manifest will be saved during a backup. Saving the manifest is important because it stores the checksums and allows the resume function to work efficiently. The actual threshold used is 1% of the backup size or `manifest-save-threshold` , whichever is greater.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024.

```
default: 1073741824
allowed: 1-1099511627776
example: manifest-save-threshold=5G
```

## Resume Option ( `-resume` )

Allow resume of failed backup.

Defines whether the resume feature is enabled. Resume can greatly reduce the amount of time required to run a backup after a previous backup of the same type has failed. It adds complexity, however, so it may be desirable to disable in environments that do not require the feature.

```
default: y
example: resume=n
```

## Start Fast Option ( `-start-fast` )

Force a checkpoint to start backup quickly.

Forces a checkpoint (by passing `y` to the `fast` parameter of `pg_start_backup()` ) so the backup begins immediately. Otherwise the backup will start after the next regular checkpoint.

This feature only works in PostgreSQL  $\geq 8.4$  .

```
default: n
example: start-fast=y
```

## Stop Auto Option ( `-stop-auto` )

Stop prior failed backup on new backup.

This will only be done if an exclusive advisory lock can be acquired to demonstrate that the prior failed backup process has really stopped.

This feature relies on `pg_is_in_backup()` so only works on PostgreSQL  $\geq 9.3$  .

The setting is disabled by default because it assumes that `pgBackRest` is the only process doing exclusive online backups. It depends on an advisory lock that only `pgBackRest` sets so it may abort other processes that do exclusive online backups. Note that `base_backup` and `pg_dump` are safe to use with this setting because they do not call `pg_start_backup()` so are not exclusive.

```
default: n
example: stop-auto=y
```

## General Options ( `general` )

The general section defines options that are common for many commands.

## Buffer Size Option ( `-buffer-size` )

Buffer size for file operations.

Set the buffer size used for copy, compress, and uncompress functions. A maximum of 3 buffers will be in use at a time per process. An additional maximum of 256K per process may be used for zlib buffers.

Size can be entered in bytes (default) or KB, MB, GB, TB, or PB where the multiplier is a power of 1024. For example, the case-insensitive value `32k` (or `32KB`) can be used instead of `32768`.

Allowed values, in bytes, are 16384 , 32768 , 65536 , 131072 , 262144 , 524288 , 1048576 , 2097152 , 4194304 , 8388608 , and 16777216 .

```
default: 4194304
example: buffer-size=32K
```

## SSH client command Option ( `-cmd-ssh` )

Path to ssh client executable.

Use a specific SSH client when an alternate is desired or the ssh executable is not in `$PATH`.

```
default: ssh
example: cmd-ssh=/usr/bin/ssh
```

## Compress Option ( `-compress` )

Use gzip file compression.

Backup files are compatible with command-line gzip tools.

```
default: y
example: compress=n
```

## Compress Level Option ( `-compress-level` )

Compression level for stored files.

Sets the zlib level to be used for file compression when `compress=y` .

```
default: 6
allowed: 0-9
example: compress-level=9
```

## Network Compress Level Option ( `-compress-level-network` )

Compression level for network transfer when `compress=n` .

Sets the zlib level to be used for protocol compression when `compress=n` and the database cluster is not on the same host as the repository. Protocol compression is used to reduce network traffic but can be disabled by setting `compress-level-network=0` . When `compress=y` the `compress-level-network` setting is ignored and `compress-level` is used instead so that the file is only compressed once. SSH compression is always disabled.

```
default: 3
allowed: 0-9
example: compress-level-network=1
```

## Database Timeout Option ( `-db-timeout` )

Database query timeout.

Sets the timeout, in seconds, for queries against the database. This includes the `pg_start_backup()` and `pg_stop_backup()` functions which can each take a substantial amount of time. Because of this the timeout should be kept high unless you know that these functions will return quickly (i.e. if you have set `startfast=y` and you know that the database cluster will not generate many WAL segments during the backup).

```
default: 1800
allowed: 0.1-604800
example: db-timeout=600
```

## Delta Option ( `-delta` )

Restore or backup using checksums.

During a restore, by default the PostgreSQL data and tablespace directories are expected to be present but empty. This option performs a delta restore using checksums.

During a backup, this option will use checksums instead of the timestamps to determine if files will be copied.

```
default: n
example: delta=y
```

## Lock Path Option ( `-lock-path` )

Path where lock files are stored.

The lock path provides a location for pgBackRest to create lock files to prevent conflicting operations from being run concurrently.

```
default: /tmp/pgbackrest
example: lock-path=/backup/db/lock
```

## Neutral Umask Option ( `-neutral-umask` )

Use a neutral umask.

Sets the umask to 0000 so modes in the repository are created in a sensible way. The default directory mode is 0750 and default file mode is 0640. The lock and log directories set the directory and file mode to 0770 and 0660 respectively.

To use the executing user's umask instead specify `neutral-umask=n` in the config file or `-no-neutral-umask` on the command line.

```
default: y
example: neutral-umask=n
```

## Process Maximum Option ( `-process-max` )

Max processes to use for compress/transfer.

Each process will perform compression and transfer to make the command run faster, but don't set `process-max` so high that it impacts database performance.

```
default: 1
allowed: 1-96
example: process-max=4
```

## Protocol Timeout Option ( `-protocol-timeout` )

Protocol timeout.

Sets the timeout, in seconds, that the local or remote process will wait for a new message to be received on the protocol layer. This prevents processes from waiting indefinitely for a message. The `protocol-timeout` option must be greater than the `db-timeout` option.

```
default: 1830
allowed: 0.1-604800
example: protocol-timeout=630
```

## Spool Path Option ( `-spool-path` )

Path where transient data is stored.

This path is used to store data for the asynchronous `archive-push` and `archive-get` command.

The asynchronous `archive-push` command writes acknowledgements into the spool path when it has successfully stored WAL in the archive (and errors on failure) so the foreground process can quickly notify PostgreSQL. Acknowledgement files are very small (zero on success and a few hundred bytes on error).

The asynchronous `archive-push` process queues WAL in the spool path so it can be provided very quickly when PostgreSQL requests it. Moving files to PostgreSQL is most efficient when the spool path is on the same filesystem as `pg_xlog` / `pg_wal`.

The data stored in the spool path is not strictly temporary since it can and should survive a reboot. However, loss of the data in the spool path is not a problem. `pgBackRest` will simply recheck each WAL segment to ensure it is safely archived for `archive-push` and rebuild the queue for `archive-get`.

The spool path is intended to be located on a local Posix-compatible filesystem, not a remote filesystem such as NFS or CIFS.

```
default: /var/spool/pgbackrest
example: spool-path=/backup/db/spool
```

## Log Options ( `log` )

The log section defines logging-related settings.

CAUTION:

Trace-level logging may expose secrets such as keys and passwords. Use with caution!

## Console Log Level Option ( `-log-level-console` )

Level for console logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: log-level-console=error
```

## File Log Level Option ( `-log-level-file` )

Level for file logging.

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: info
```

```
example: log-level-file=debug
```

## Std Error Log Level Option ( `-log-level-stderr` )

Level for stderr logging.

Specifies which log levels will output to stderr rather than stdout (specified by `log-level-console` ). The timestamp and process will not be output to stderr .

The following log levels are supported:

- off - No logging at all (not recommended)
- error - Log only errors
- warn - Log warnings and errors
- info - Log info, warnings, and errors
- detail - Log detail, info, warnings, and errors
- debug - Log debug, detail, info, warnings, and errors
- trace - Log trace (very verbose debugging), debug, info, warnings, and errors

```
default: warn
```

```
example: log-level-stderr=error
```



## Log Path Option ( `-log-path` )

Path where log files are stored.

The log path provides a location for pgBackRest to store log files. Note that if `log-level-file=off` then no log path is required.

```
default: /var/log/pgbackrest
example: log-path=/backup/db/log
```

## Log Subprocesses Option ( `-log-subprocess` )

Enable logging in subprocesses.

Enable file logging for any subprocesses created by this process using the log level specified by `log-level-file` .

```
default: n
example: log-subprocess=y
```

## Log Timestamp Option ( `-log-timestamp` )

Enable timestamp in logging.

Enables the timestamp in console and file logging. This option is disabled in special situations such as generating documentation.

```
default: y
example: log-timestamp=n
```

## Repository Options ( `repository` )

The repository section defines options used to configure the repository.

**Indexing** : All `repo-` options are indexed to allow for configuring multiple repositories, though only a single repository is currently supported. For example, the repository is configured with the `repo1-path` , `repo1-host` , etc. options.

The `repo-retention-*` options define how long backups will be retained. Expiration only occurs when the number of complete backups exceeds the allowed retention. In other words, if `retention-full` is set to 2, then there must be 3 complete backups before the oldest will be expired. Make sure you always have enough space for `retention + 1` backups.

## Repository Cipher Passphrase Option ( `-repo-cipher-pass` )

Repository cipher passphrase.

Passphrase used to encrypt/decrypt files of the repository.

```
example: repo1-cipher-pass=zWaf6XtpjIVZC5444yXB+cgFDF17MxGlgkZSaoPvTGirhPygu4j0K0Xf9L04vjf0
```

## Repository Cipher Type Option ( `-repo-cipher-type` )

Cipher used to encrypt the repository.

The following repository types are supported:

- `none` - The repository is not encrypted
- `aes-256-cbc` - Advanced Encryption Standard with 256 bit key length

Note that encryption is always performed client-side even if the repository type (e.g. `S3`) supports encryption.

```
default: none
example: repo1-cipher-type=aes-256-cbc
```

## Repository Hardlink Option ( `-repo-hardlink` )

Hardlink files between backups in the repository.

Enable hard-linking of files in differential and incremental backups to their full backups. This gives the appearance that each backup is a full backup at the file-system level. Be careful, though, because modifying files that are hard-linked can affect all the backups in the set.

```
default: n
example: repo1-hardlink=y
```

Deprecated Name: `hardlink`

## Repository Host Option ( `-repo-host` )

Repository host when operating remotely via SSH.

Make sure that trusted SSH authentication is configured between the PostgreSQL host and the repository host.

When backing up and archiving to a locally mounted filesystem this setting is not required.

```
example: repo1-host=repo1.domain.com
```

Deprecated Name: `backup-host`

## Repository Host Command Option ( `-repo-host-cmd` )

pgBackRest exe path on the repository host.

Required only if the path to `pgbackrest` is different on the local and repository hosts. If not defined, the repository host exe path will be set the same as the local exe path.

```
example: repo1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: `backup-cmd`

## Repository Host Configuration Option ( `-repo-host-config` )

pgBackRest repository host configuration file.

Sets the location of the configuration file on the repository host. This is only required if the repository host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
example: repo1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: `backup-config`

## Repository Host Configuration Include Path Option ( `-repo-host-config-include-path` )

pgBackRest repository host configuration include path.

Sets the location of the configuration include path on the repository host. This is only required if the repository host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
example: repo1-host-config-include-path=/conf/pgbackrest/conf.d
```

## Repository Host Configuration Path Option ( `-repo-host-config-path` )

pgBackRest repository host configuration path.

Sets the location of the configuration path on the repository host. This is only required if the repository host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: repo1-host-config-path=/conf/pgbackrest
```

## Repository Host Port Option ( `-repo-host-port` )

Repository host port when repo-host is set.

Use this option to specify a non-default port for the repository host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: repo1-host-port=25
```

Deprecated Name: backup-ssh-port

## Repository Host User Option ( `-repo-host-user` )

Repository host user when repo-host is set.

Defines the user that will be used for operations on the repository host. Preferably this is not the postgres user but rather some other user like pgbackrest . If PostgreSQL runs on the repository host the postgres user can be placed in the pgbackrest group so it has read permissions on the repository without being able to damage the contents accidentally.

```
default: pgbackrest
example: repo1-host-user=repo-user
```

Deprecated Name: backup-user

## Repository Path Option ( `-repo-path` )

Path where backups and archive are stored.

The repository is where pgBackRest stores backups and archives WAL segments.

It may be difficult to estimate in advance how much space you'll need. The best thing to do is take some backups then record the size of different types of backups (full/incr/diff) and measure the amount of WAL generated per day. This will give you a general idea of how much space you'll need, though of course requirements will likely change over time as your database evolves.

```
default: /var/lib/pgbackrest
example: repo1-path=/backup/db/backrest
```

## Archive Retention Option ( `-repo-retention-archive` )

Number of backups worth of continuous WAL to retain.

NOTE:

WAL segments required to make a backup consistent are always retained until the backup is expired regardless of how this option is configured.

If this value is not set, then the archive to expire will default to the repo-retention-full (or repo-retention-diff ) value corresponding to the repo-retention-archive-type if set to full (or diff ). This will ensure that WAL is only expired for backups that are already expired.

This option must be set if repo-retention-archive-type is set to incr . If disk space is at a premium, then this setting, in conjunction with repo-retention-archive-type , can be used to aggressively expire WAL segments. However, doing so negates the ability to perform PITR from the backups with expired WAL and is therefore **not** recommended.

```
allowed: 1-9999999
example: repo1-retention-archive=2
```

Deprecated Name: retention-archive

## Archive Retention Type Option ( `-repo-retention-archive-type` )

Backup type for WAL retention.

If set to full pgBackRest will keep archive logs for the number of full backups defined by repo-retention-archive . If set to diff (differential) pgBackRest will keep archive logs for the number of full and differential backups defined by repo-retention-archive , meaning if the last backup taken was a full backup, it will be counted as a differential for the purpose of repo-retention. If set to incr (incremental) pgBackRest will keep archive logs for the number of full, differential, and incremental backups defined by repo-retention-archive . It is recommended that this setting not be changed from the default which will only expire WAL in conjunction with expiring full backups.

```
default: full
example: repo1-retention-archive-type=diff
```

Deprecated Name: retention-archive-type

### Differential Retention Option ( `-repo-retention-diff` )

Number of differential backups to retain.

When a differential backup expires, all incremental backups associated with the differential backup will also expire. When not defined all differential backups will be kept until the full backups they depend on expire.

```
allowed: 1-9999999
example: repo1-retention-diff=3
```

Deprecated Name: retention-diff

### Full Retention Option ( `-repo-retention-full` )

Number of full backups to retain.

When a full backup expires, all differential and incremental backups associated with the full backup will also expire. When the option is not defined a warning will be issued. If indefinite retention is desired then set the option to the max value.

```
allowed: 1-9999999
example: repo1-retention-full=2
```

Deprecated Name: retention-full

### S3 Repository Bucket Option ( `-repo-s3-bucket` )

S3 repository bucket.

S3 bucket used to store the repository.

pgBackRest repositories can be stored in the bucket root by setting `repo-path=/` but it is usually best to specify a prefix, such as `/repo`, so logs and other AWS generated content can also be stored in the bucket.

```
example: repo1-s3-bucket=pg-backup
```

### S3 SSL CA File Option ( `-repo-s3-ca-file` )

S3 SSL CA File.

Use a CA file other than the system default.

```
example: repo1-s3-ca-file=/etc/pki/tls/certs/ca-bundle.crt
```

### S3 SSL CA Path Option ( `-repo-s3-ca-path` )

S3 SSL CA Path.

Use a CA path other than the system default.

```
example: repo1-s3-ca-path=/etc/pki/tls/certs
```

### S3 Repository Endpoint Option ( `-repo-s3-endpoint` )

S3 repository endpoint.

The AWS end point should be valid for the selected region.

```
example: repo1-s3-endpoint=s3.amazonaws.com
```

### **S3 Repository Host Option ( `-repo-s3-host` )**

S3 repository host.

Connect to a host other than the end point. This is typically used for testing.

```
example: repo1-s3-host=127.0.0.1
```

### **S3 Repository Access Key Option ( `-repo-s3-key` )**

S3 repository access key.

AWS key used to access this bucket.

```
example: repo1-s3-key=AKIAIOSFODNN7EXAMPLE
```

### **S3 Repository Secret Access Key Option ( `-repo-s3-key-secret` )**

S3 repository secret access key.

AWS secret key used to access this bucket.

```
example: repo1-s3-key-secret=wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
```

### **S3 Repository Region Option ( `-repo-s3-region` )**

S3 repository region.

The AWS region where the bucket was created.

```
example: repo1-s3-region=us-east-1
```

### **S3 Repository Security Token Option ( `-repo-s3-token` )**

S3 repository security token.

AWS security token used with temporary credentials.

```
example: repo1-s3-token=AQoDYXdzEPT//////////wEXAMPLEetc764bNrC9SAPBSM22 ...
```

### **S3 Repository Verify SSL Option ( `-repo-s3-verify-ssl` )**

Verify S3 server certificate.

Disables verification of the S3 server certificate. This should only be used for testing or other scenarios where a certificate has been self-signed.

```
default: y
```

```
example: repo1-s3-verify-ssl=n
```

### **Repository Type Option ( `-repo-type` )**

Type of storage used for the repository.

The following repository types are supported:

- cifs - Like posix , but disables links and directory fsyncs
- posix - Posix-compliant file systems
- s3 - AWS Simple Storage Service

```
default: posix
```

```
example: repo1-type=cifs
```

## Restore Options ( `restore` )

The restore section defines settings used for restoring backups.

### Include Database Option ( `-db-include` )

Restore only specified databases.

This feature allows only selected databases to be restored. Databases not specifically included will be restored as sparse, zeroed files to save space but still allow PostgreSQL to perform recovery. After recovery the databases that were not included will not be accessible but can be removed with the drop database command.

NOTE:

built-in databases ( `template0` , `template1` , and `postgres` ) are always restored.

The `-db-include` option can be passed multiple times to specify more than one database to include.

```
example: db-include=db_main
```

### Link All Option ( `-link-all` )

Restore all symlinks.

By default symlinked directories and files are restored as normal directories and files in `$PGDATA`. This is because it may not be safe to restore symlinks to their original destinations on a system other than where the original backup was performed. This option restores all the symlinks just as they were on the original system where the backup was performed.

```
default: n
```

```
example: link-all=y
```

### Link Map Option ( `-link-map` )

Modify the destination of a symlink.

Allows the destination file or path of a symlink to be changed on restore. This is useful for restoring to systems that have a different storage layout than the original system where the backup was generated.

```
example: link-map=pg_xlog=/data/xlog
```

### Recovery Option Option ( `-recovery-option` )

Set an option in `recovery.conf` .

See <http://www.postgresql.org/docs/X.X/static/recovery-config.html> for details on `recovery.conf` options (replace `X.X` with your PostgreSQL version). This option can be used multiple times.

NOTE:

The `restore_command` option will be automatically generated but can be overridden with this option. Be careful about specifying your own `restore_command` as `pgBackRest` is designed to handle this for you. Target Recovery options (`recovery_target_name`, `recovery_target_time`, etc.) are generated automatically by `pgBackRest` and should not be set with this option.

Since `pgBackRest` does not start PostgreSQL after writing the `recovery.conf` file, it is always possible to edit/check `recovery.conf` before manually restarting.

```
example: recovery-option=primary_conninfo=db.mydomain.com
```

### Tablespace Map Option ( `-tablespace-map` )

Restore a tablespace into the specified directory.

Moves a tablespace to a new location during the restore. This is useful when tablespace locations are not the same on a replica, or an upgraded system has different mount points.

Since PostgreSQL 9.2 tablespace locations are not stored in `pg_tablespace` so moving tablespaces can be done with impunity. However, moving a tablespace to the `data_directory` is not recommended and may cause problems. For more information on moving tablespaces <http://www.databasesoup.com/2013/11/moving-tablespaces.html> is a good resource.

```
example: tablespace-map=ts_01=/db/ts_01
```

## Map All Tablespaces Option ( `-tablespace-map-all` )

Restore all tablespaces into the specified directory.

By default tablespaces are restored into their original locations and while this behavior can be modified by with the `tablespace-map` open it is sometime preferable to remap all tablespaces to a new directory all at once. This is particularly useful for development or staging systems that may not have the same storage layout as the original system where the backup was generated.

The path specified will be the parent path used to create all the tablespaces in the backup.

```
example: tablespace-map-all=/data/tablespace
```

## Stanza Options ( `stanza` )

A stanza defines the backup configuration for a specific PostgreSQL database cluster. The stanza section must define the database cluster path and host/user if the database cluster is remote. Also, any global configuration sections can be overridden to define stanza-specific settings.

**Indexing** : All `pg-` options are indexed to allow for configuring multiple PostgreSQL hosts. For example, a single master is configured with the `pg1-path` , `pg1-host` , etc. options. If a standby is configured then index the `pg-` options as `pg2-` (e.g. `pg2-host`, `pg2-path`, etc).

## PostgreSQL Host Option ( `-pg-host` )

PostgreSQL host for operating remotely via SSH.

Used for backups where the PostgreSQL host is different from the repository host.

```
example: pg1-host=db.domain.com
```

Deprecated Name: `db-host`

## PostgreSQL Host Command Option ( `-pg-host-cmd` )

`pgBackRest` exe path on the PostgreSQL host.

Required only if the path to `pgbackrest` is different on the local and PostgreSQL hosts. If not defined, the database host exe path will be set the same as the local exe path.

```
example: pg1-host-cmd=/usr/lib/backrest/bin/pgbackrest
```

Deprecated Name: `db-cmd`

## PostgreSQL Host Configuration Option ( `-pg-host-config` )

`pgBackRest` database host configuration file.

Sets the location of the configuration file on the PostgreSQL host. This is only required if the PostgreSQL host configuration file is in a different location than the local configuration file.

```
default: /etc/pgbackrest/pgbackrest.conf
```

```
example: pg1-host-config=/conf/pgbackrest/pgbackrest.conf
```

Deprecated Name: `db-config`

## PostgreSQL Host Configuration Include Path Option ( `-pg-host-config-include-path` )

`pgBackRest` database host configuration include path.

Sets the location of the configuration include path on the PostgreSQL host. This is only required if the PostgreSQL host configuration include path is in a different location than the local configuration include path.

```
default: /etc/pgbackrest/conf.d
```

```
example: pg1-host-config-include-path=/conf/pgbackrest/conf.d
```

## PostgreSQL Host Configuration Path Option ( `-pg-host-config-path` )

pgBackRest database host configuration path.

Sets the location of the configuration path on the PostgreSQL host. This is only required if the PostgreSQL host configuration path is in a different location than the local configuration path.

```
default: /etc/pgbackrest
example: pg1-host-config-path=/conf/pgbackrest
```

## PostgreSQL Host Port Option ( `-pg-host-port` )

PostgreSQL host port when pg-host is set.

Use this option to specify a non-default port for the PostgreSQL host protocol. Currently only SSH is supported

```
allowed: 0-65535
example: pg1-host-port=25
```

Deprecated Name: db-ssh-port

## PostgreSQL Host User Option ( `-pg-host-user` )

PostgreSQL host logon user when pg-host is set.

This user will also own the remote pgBackRest process and will initiate connections to PostgreSQL . For this to work correctly the user should be the PostgreSQL database cluster owner which is generally postgres , the default.

```
default: postgres
example: pg1-host-user=db_owner
```

Deprecated Name: db-user

## PostgreSQL Path Option ( `-pg-path` )

PostgreSQL data directory.

This should be the same as the data\_directory setting in postgresql.conf . Even though this value can be read from postgresql.conf or PostgreSQL it is prudent to set it in case those resources are not available during a restore or offline backup scenario.

The pg-path option is tested against the value reported by PostgreSQL on every online backup so it should always be current.

```
example: pg1-path=/data/db
```

Deprecated Name: db-path

## PostgreSQL Port Option ( `-pg-port` )

PostgreSQL port.

Port that PostgreSQL is running on. This usually does not need to be specified as most PostgreSQL clusters run on the default port.

```
default: 5432
allowed: 0-65535
example: pg1-port=6543
```

Deprecated Name: db-port

## PostgreSQL Socket Path Option ( `-pg-socket-path` )

PostgreSQL unix socket path.

The unix socket directory that was specified when PostgreSQL was started. pgBackRest will automatically look in the standard location for your OS so there is usually no need to specify this setting unless the socket directory was explicitly modified with the unix\_socket\_directory setting in postgresql.conf .

```
allowed: 0-65535
example: pg1-socket-path=/var/run/postgresql
```



## Introduction

pgBackRest release numbers consist of two parts, major and minor. A major release *may* break compatibility with the prior major release, but v2 releases are fully compatible with v1 repositories and will accept all v1 options. Minor releases can include bug fixes and features but do not change the repository format and strive to avoid changing options and naming.

Documentation for the v1 release can be found [here](#) .

The notes for a release may also contain “Additional Notes” but changes in this section are only to documentation or the test suite and have no direct impact on the pgBackRest codebase.

## Current Stable Release

### v2.11 Release Notes

C Implementation of Archive Get

Released March 11, 2019

#### Bug Fixes:

- Fix possible truncated WAL segments when an error occurs mid-write. ( *Reported by blogh.* )
- Fix info command missing WAL min/max when stanza specified. ( *Fixed by Stefan Fercot.* )
- Fix non-compliant JSON for options passed from C to Perl. ( *Reported by Leo Khomenko.* )

#### Improvements:

- The archive-get command is implemented entirely in C.
- Enable socket keep-alive on older Perl versions. ( *Contributed by Marc Cousin.* )
- Error when parameters are passed to a command that does not accept parameters. ( *Suggested by Jason O'Donnell.* )
- Add hints when unable to find a WAL segment in the archive. ( *Suggested by Hans-Jürgen Schöning.* )
- Improve error when hostname cannot be found in a certificate. ( *Suggested by James Badger.* )
- Add additional options to backup.manifest for debugging purposes. ( *Contributed by blogh.* )

Additional Notes

#### Documentation Improvements:

- Update default documentation version to PostgreSQL 10.

## Stable Releases

### v2.10 Release Notes

Bug Fixes

Released February 9, 2019

#### Bug Fixes:

- Add unimplemented S3 driver method required for archive-get . ( *Reported by mibvio.* )
- Fix check for improperly configured pg-path . ( *Reported by James Chanco Jr.* )

## v2.09 Release Notes

Minor Improvements and Bug Fixes

Released January 30, 2019

### Bug Fixes:

- Fix issue with multiple async status files causing a hard error. ( *Reported by Vidhya Gurumoorthi, Joe Ayers, Douglas J Hunley.* )

### Improvements:

- The info command is implemented entirely in C.
- Simplify info command text message when no stanzas are present by replacing the repository path with “the repository” .
- Add `_DARWIN_C_SOURCE` flag to Makefile for MacOS builds. ( *Contributed by Douglas J Hunley.* )
- Update address lookup in C TLS client to use modern methods. ( *Suggested by Bruno Friedmann.* )
- Include Posix-compliant header for `strcasecmp()` and `fd_set` . ( *Suggested by ucando.* )

Additional Notes

### Documentation Bug Fixes:

- Fix hard-coded repository path. ( *Reported by Heath Lord.* )

### Documentation Improvements:

- Clarify that encryption is always performed client-side. ( *Suggested by Bruce Burdick.* )
- Add examples for building a documentation host.
- Allow if in manifest variables, lists, and list items.

## v2.08 Release Notes

Minor Improvements and Bug Fixes

Released January 2, 2019

### Bug Fixes:

- Remove request for S3 object info directly after putting it. ( *Reported by Matt Kunkel.* )
- Correct archive-get-queue-max to be size type. ( *Reported by Ronan Dunklau.* )
- Add error message when current user uid / gid does not map to a name. ( *Reported by Camilo Aguilar.* )
- Error when `-target-action=shutdown` specified for PostgreSQL < 9.5.

### Improvements:

- Set TCP keepalives on S3 connections. ( *Suggested by Ronan Dunklau.* )
- Reorder info command text output so most recent backup is output last. ( *Contributed by Cynthia Shang. Suggested by Ryan Lambert.* )
- Change file ownership only when required.
- Redact authentication header when throwing S3 errors. ( *Suggested by Brad Nicholson.* )

Additional Notes

### Documentation Improvements:

- Clarify when target-action is effective and PostgreSQL version support. ( *Suggested by Keith Fiske.* )
- Clarify that region/endpoint must be configured correctly for the bucket. ( *Suggested by Pritam Barhate.* )
- Add documentation for building the documentation.

## v2.07 Release Notes

Automatic Backup Checksum Delta

Released November 16, 2018

### Bug Fixes:

- Fix issue with archive-push-queue-max not being honored on connection error. ( *Reported by Lardière Sébastien.* )
- Fix static WAL segment size used to determine if archive-push-queue-max has been exceeded.
- Fix error after log file open failure when processing should continue. ( *Reported by vthriller.* )

### Features:

- Automatically enable backup checksum delta when anomalies (e.g. timeline switch) are detected. ( *Contributed by Cynthia Shang.* )

### Improvements:

- Retry all S3 5xx errors rather than just 500 internal errors. ( *Suggested by Craig A. James.* )

## v2.06 Release Notes

Checksum Delta Backup and PostgreSQL 11 Support

Released October 15, 2018

### Bug Fixes:

- Fix missing URI encoding in S3 driver. ( *Reported by Dan Farrell.* )
- Fix incorrect error message for duplicate options in configuration files. ( *Reported by Jesper St John.* )
- Fix incorrectly reported error return in info logging. A return code of 1 from the archive-get was being logged as an error message at info level but otherwise worked correctly.

### Features:

- Add checksum delta for incremental backups which uses checksums rather than timestamps to determine if files have changed. ( *Contributed by Cynthia Shang.* )
- PostgreSQL 11 support, including configurable WAL segment size.

### Improvements:

- Ignore all files in a linked tablespace directory except the subdirectory for the current version of PostgreSQL . Previously an error would be generated if other files were present and not owned by the PostgreSQL user.
- Improve info command to display the stanza cipher type. ( *Contributed by Cynthia Shang. Suggested by Douglas J Hunley.* )
- Improve support for special characters in filenames.
- Allow delta option to be specified in the pgBackRest configuration file. ( *Contributed by Cynthia Shang.* )

Additional Notes

### Documentation Improvements:

- Use command in authorized\_hosts to improve SSH security. ( *Suggested by Stephen Frost, Magnus Hagander.* )
- List allowable values for the buffer-size option in the configuration reference. ( *Contributed by Cynthia Shang. Suggested by Stéphane Schildknecht.* )

## v2.05 Release Notes

Environment Variable Options and Exclude Temporary/Unlogged Relations

Released August 31, 2018

### Bug Fixes:

- Fix issue where *relative* links in \$PGDATA could be stored in the backup with the wrong path. This issue did not affect absolute links and relative tablespace links were caught by other checks. ( *Reported by Cynthia Shang.* )
- Remove incompletely implemented online option from the check command. Offline operation runs counter to the purpose of this command, which is to check if archiving and backups are working correctly. ( *Reported by Jason O'Donnell.* )
- Fix issue where errors raised in C were not logged when called from Perl. pgBackRest properly terminated with the correct error code but lacked an error message to aid in debugging. ( *Reported by Douglas J Hunley.* )
- Fix issue when a boolean option (e.g. delta ) was specified more than once. ( *Reported by Yogesh Sharma.* )

### Features:

- Allow any option to be set in an environment variable. This includes options that previously could only be specified on the command line, e.g. stanza , and secret options that could not be specified on the command-line, e.g. repo1-s3-key-secret .
- Exclude temporary and unlogged relation (table/index) files from backup. Implemented using the same logic as the patches adding this feature to PostgreSQL , [8694cc96](#) and [920a5e50](#) . Temporary relation exclusion is enabled in PostgreSQL 9.0 . Unlogged relation exclusion is enabled in PostgreSQL 9.1 , where the feature was introduced. ( *Contributed by Cynthia Shang.* )
- Allow arbitrary directories and/or files to be excluded from a backup. Misuse of this feature can lead to inconsistent backups so read the `-exclude` documentation carefully before using. ( *Reviewed by Cynthia Shang.* )
- Add log-subprocess option to allow file logging for local and remote subprocesses.
- PostgreSQL 11 Beta 3 support.

### Improvements:

- Allow zero-size files in backup manifest to reference a prior manifest regardless of timestamp delta. ( *Contributed by Cynthia Shang.* )
- Improve asynchronous archive-get / archive-push performance by directly checking status files. ( *Contributed by Stephen Frost.* )
- Improve error message when a command is missing the stanza option. ( *Suggested by Sarah Conway.* )

Additional Notes

### Documentation Bug Fixes:

- Fix invalid log level in log-path option reference. ( *Reported by Camilo Aguilar.* )

### Documentation Improvements:

- Stop trying to arrange contributors in release.xml by last/first name. Contributor names have always been presented in the release notes exactly as given, but we tried to assign internal IDs based on last/first name which can be hard to determine and ultimately doesn't make sense. Inspired by Christophe's PostgresOpen 2017 talk, "Human Beings Do Not Have a Primary Key" . ( *Suggested by Christophe Pettus.* )

### Test Suite Improvements:

- Error if LibC build is performed outside the test environment. LibC is no longer required for production builds.

## v2.04 Release Notes

Critical Bug Fix for Backup Resume

Released July 5, 2018

**IMPORTANT NOTE** : This release fixes a critical bug in the backup resume feature. All resumed backups prior to this release should be considered inconsistent. A backup will be resumed after a prior backup fails, unless `resume=n` has been specified. A resumed backup can be identified by checking the backup log for the message "aborted backup of same type exists, will be cleaned to remove invalid files and resumed" . If the message exists, do not use this backup or any backup in the same set for a restore and check the restore logs to see if a resumed backup was restored. If so, there may be inconsistent data in the cluster.

### Bug Fixes:

- Fix critical bug in resume that resulted in inconsistent backups. A regression in v0.82 removed the timestamp comparison when deciding which files from the aborted backup to keep on resume. See note above for more details. ( *Reported by David Youatt, Yogesh Sharma, Stephen Frost.* )
- Fix error in selective restore when only one user database exists in the cluster. ( *Fixed by Cynthia Shang. Reported by Nj Baliyan.* )
- Fix non-compliant ISO-8601 timestamp format in S3 authorization headers. AWS and some gateways were tolerant of space rather than zero-padded hours while others were not. ( *Fixed by Andrew Schwartz.* )

#### Features:

- PostgreSQL 11 Beta 2 support.

#### Improvements:

- Improve the HTTP client to set content-length to 0 when not specified by the server. S3 (and gateways) always set content-length or transfer-encoding but HTTP 1.1 does not require it and proxies (e.g. HAProxy ) may not include either. ( *Suggested by Adam K. Sumner.* )
- Set search\_path = 'pg\_catalog' on PostgreSQL connections. ( *Suggested by Stephen Frost.* )

#### Additional Notes

#### Documentation Improvements:

- Create a new section to describe building pgBackRest and build on a separate host.
- Add sample S3 policy to restrict bucket privileges. ( *Suggested by Douglas J Hunley, Jason O'Donnell.* )

### v2.03 Release Notes

#### Single Executable to Deploy

Released May 22, 2018

#### Bug Fixes:

- Fix potential buffer overrun in error message handling. ( *Reported by Lætitia.* )
- Fix archive write lock being taken for the synchronous archive-get command. ( *Reported by uspen.* )

#### Improvements:

- Embed exported C functions and Perl modules directly into the pgBackRest executable.
- Use time\_t instead of \_\_time\_t for better portability. ( *Suggested by Nick Floersch.* )
- Print total runtime in milliseconds at command end.

### v2.02 Release Notes

#### Parallel Asynchronous Archive Get and Configuration Includes

Released May 6, 2018

#### Bug Fixes:

- Fix directory syncs running recursively when only the specified directory should be synced. ( *Reported by Craig A. James.* )
- Fix archive-copy throwing “path not found” error for incr/diff backups. ( *Reported by yummyliu, Vitaliy Kukharik.* )
- Fix failure in manifest build when two or more files in PGDATA are linked to the same directory. ( *Reported by Vitaliy Kukharik.* )
- Fix delta restore failing when a linked file is missing.
- Fix rendering of key/value and list options in help. ( *Reported by Clinton Adams.* )

#### Features:

- Add asynchronous, parallel archive-get . This feature maintains a queue of WAL segments to help reduce latency when PostgreSQL requests a WAL segment with restore\_command .
- Add support for additional pgBackRest configuration files in the directory specified by the -config-include-path option. Add -config-path option for overriding the default base path of the -config and -config-include-path option. ( *Contributed by Cynthia Shang.* )

- Add `repo-s3-token` option to allow temporary credentials tokens to be configured. `pgBackRest` currently has no way to request new credentials so the entire command (e.g. `backup` , `restore` ) must complete before the credentials expire. ( *Contributed by Yogesh Sharma.* )

### Improvements:

- Update the `archive-push-queue-max` , `manifest-save-threshold` , and `buffer-size` options to accept values in KB , MB , GB , TB , or PB where the multiplier is a power of 1024 . ( *Contributed by Cynthia Shang.* )
- Make backup/restore path sync more efficient. Scanning the entire directory can be very expensive if there are a lot of small tables. The backup manifest contains the path list so use it to perform syncs instead of scanning the backup/restore path.
- Show command parameters as well as command options in initial info log message.
- Rename `archive-queue-max` option to `archive-push-queue-max` to avoid confusion with the new `archive-get-queue-max` option. The old option name will continue to be accepted.

### Additional Notes

### Documentation Bug Fixes:

- Update docs with 32-bit support and caveats. 32-bit support was added in v1.26 . ( *Reported by Viorel Tabara.* )

### Documentation Improvements:

- Add monitoring examples using PostgreSQL and `jq` . ( *Suggested by Stephen Frost, Brian Faherty.* )
- Add example of command section usage to archiving configuration. ( *Suggested by Christophe Courtois.* )
- Remove documentation describing `info -output=json` as experimental.
- Update out-of-date description for the `spool-path` option.

### Test Suite Features:

- Use `lcov` for C unit test coverage reporting. Switch from `Devel::Cover` because it would not report on branch coverage for reports converted from `gcov` . Incomplete branch coverage for a module now generates an error. Coverage of unit tests is not displayed in the report unless they are incomplete for either statement or branch coverage.

## v2.01 Release Notes

### Minor Bug Fixes and Improvements

Released March 19, 2018

### Bug Fixes:

- Fix `-target-action` and `-recovery-option` options being reported as invalid when restoring with `-type=immediate` . ( *Reported by Brad Nicholson.* )
- Immediately error when a secure option (e.g. `repo1-s3-key` ) is passed on the command line. Since `pgBackRest` would not pass secure options on to sub-processes an obscure error was thrown. The new error is much clearer and provides hints about how to fix the problem. Update command documentation to omit secure options that cannot be specified on the command-line. ( *Reported by Brad Nicholson.* )
- Fix issue passing `-no-config` to embedded Perl. ( *Reported by Ibrahim Edib Kokdemir.* )
- Fix issue where specifying `log-level-stderr > warn` would cause a local / remote process to error on exit due to output found on `stderr` when none was expected. The max value for a local / remote process is now error since there is no reason for these processes to emit warnings. ( *Reported by Clinton Adams.* )
- Fix manifest test in the `check` command when tablespaces are present. ( *Fixed by Cynthia Shang. Reported by Thomas Flatley.* )

### Improvements:

- Error when multiple arguments are set in the config file for an option that does not accept multiple arguments. ( *Contributed by Cynthia Shang.* )
- Remove extraneous `sudo` commands from `src/Makefile` . ( *Contributed by Adrian Vondendriesch.* )

### Additional Notes

### Documentation Improvements:

- Show index in examples for indexed options, i.e. `repo-*` , `pg-*` . ( *Suggested by Stephen Frost.* )
- Simplify table of contents on command page by only listing commands. ( *Suggested by Stephen Frost.* )
- Remove references to the C library being optional.

### Test Suite Features:

- Add CentOS/RHEL package builds.
- Use clang for static code analysis. Nothing found initially except for some functions that should have been marked `__noreturn__` .

### v2.00 Release Notes

#### Performance Improvements for Archive Push

Released February 23, 2018

#### Features:

- The `archive-push` command is now partially coded in C which allows the PostgreSQL `archive_command` to run significantly faster when processing status messages from the asynchronous archive process. ( *Reviewed by Cynthia Shang.* )

#### Improvements:

- Improve `check` command to verify that the backup manifest can be built. ( *Contributed by Cynthia Shang.* )
- Improve performance of HTTPS client. Buffering now takes the pending bytes on the socket into account (when present) rather than relying entirely on `select()` . In some instances the final bytes would not be flushed until the connection was closed.
- Improve S3 delete performance. The constant `S3_BATCH_MAX` had been replaced with a hard-coded value of 2, probably during testing.
- Allow any non-command-line option to be reset to default on the command-line. This allows options in `pgbackrest.conf` to be reset to default which reduces the need to write new configuration files for specific needs.
- The C library is now required. This eliminates conditional loading and eases development of new library features.
- The `pgbackrest` executable is now a C binary instead of Perl. This allows certain time-critical commands (like `async archive-push` ) to run more quickly.
- Rename `db-*` options to `pg-*` and `backup-*` options to `repo-*` to improve consistency. `repo-*` options are now indexed although currently only one is allowed.

#### Additional Notes

#### Documentation Features:

- All clusters in the documentation are initialized with checksums.

#### Documentation Improvements:

- List deprecated option names in documentation and command-line help.
- Clarify that S3 buckets must be created by the user. ( *Suggested by David Youatt.* )

### v1.29 Release Notes

#### Critical Bug Fix for Backup Resume

Released July 5, 2018

**IMPORTANT NOTE** : This release fixes a critical bug in the backup resume feature. All resumed backups prior to this release should be considered inconsistent. A backup will be resumed after a prior backup fails, unless `resume=n` has been specified. A resumed backup can be identified by checking the backup log for the message “aborted backup of same type exists, will be cleaned to remove invalid files and resumed” . If the message exists, do not use this backup or any backup in the same set for a restore and check the restore logs to see if a resumed backup was restored. If so, there may be inconsistent data in the cluster.

#### Bug Fixes:

- Fix critical bug in resume that resulted in inconsistent backups. A regression in v0.82 removed the timestamp comparison when deciding which files from the aborted backup to keep on resume. See note above for more details. ( *Reported by David Youatt, Yogesh Sharma, Stephen Frost.* )

- Fix non-compliant ISO-8601 timestamp format in S3 authorization headers. AWS and some gateways were tolerant of space rather than zero-padded hours while others were not. ( *Fixed by Andrew Schwartz.* )
- Fix directory syncs running recursively when only the specified directory should be synced. ( *Reported by Craig A. James.* )
- Fix `-target-action` and `-recovery-option` options being reported as invalid when restoring with `-type=immediate` . ( *Reported by Brad Nicholson.* )
- Fix archive-copy throwing “path not found” error for incr/diff backups. ( *Reported by yummyliu, Vitaliy Kukharik.* )
- Fix failure in manifest build when two or more files in PGDATA are linked to the same directory. ( *Reported by Vitaliy Kukharik.* )
- Fix delta restore failing when a linked file was missing.
- Fix error in selective restore when only one user database exists in the cluster. ( *Fixed by Cynthia Shang. Reported by Nj Baliyan.* )

### Improvements:

- Improve the HTTP client to set content-length to 0 when not specified by the server. S3 (and gateways) always set content-length or transfer-encoding but HTTP 1.1 does not require it and proxies (e.g. HAProxy ) may not include either. ( *Suggested by Adam K. Sumner.* )
- Improve performance of HTTPS client. Buffering now takes the pending bytes on the socket into account (when present) rather than relying entirely on `select()` . In some instances the final bytes would not be flushed until the connection was closed.
- Improve S3 delete performance. The constant `S3_BATCH_MAX` had been replaced with a hard-coded value of 2, probably during testing.
- Make backup/restore path sync more efficient. Scanning the entire directory can be very expensive if there are a lot of small tables. The backup manifest contains the path list so use it to perform syncs instead of scanning the backup/restore path. Remove recursive path sync functionality since it is no longer used.

### Additional Notes

### Documentation Bug Fixes:

- Update docs with 32-bit support and caveats. 32-bit support was added in v1.26 . ( *Reported by Viorel Tabara.* )

### Documentation Improvements:

- Clarify that S3 buckets must be created by the user. ( *Suggested by David Youatt.* )
- Update out-of-date description for the `spool-path` option.

## v1.28 Release Notes

### Stanza Delete

Released February 1, 2018

### Bug Fixes:

- Fixed inability to restore a single database contained in a tablespace using `-db-include`. ( *Fixed by Cynthia Shang. Reported by Chiranjeevi Ravilla.* )
- Ensure latest db-id is selected on when matching `archive.info` to `backup.info` . This provides correct matching in the event there are system-id and db-version duplicates (e.g. after reverting a `pg_upgrade` ). ( *Fixed by Cynthia Shang. Reported by Adam K. Sumner.* )
- Fixed overly chatty error message when reporting an invalid command. ( *Reported by Jason O'Donnell.* )

### Features:

- Add `stanza-delete` command to cleanup unused stanzas. ( *Contributed by Cynthia Shang. Suggested by Magnus Hagander.* )

### Improvements:

- Improve `stanza-create` command so that it does not error when the stanza already exists. ( *Contributed by Cynthia Shang.* )

### Additional Notes

### Documentation Improvements:

- Update `stanza-create -force` documentation to urge caution when using. ( *Suggested by Jason O'Donnell.* )



## v1.27 Release Notes

Bug Fixes and Documentation

Released December 19, 2017

### Bug Fixes:

- Fixed an issue that suppressed locality errors for backup and restore . When a backup host is present, backups should only be allowed on the backup host and restores should only be allowed on the database host unless an alternate configuration is created that ignores the remote host. ( *Reported by Lardière Sébastien.* )
- Fixed an issue where WAL was not expired on PostgreSQL 10. This was caused by a faulty regex that expected all PostgreSQL major versions to be X.X. ( *Reported by Adam Brusselback.* )
- Fixed an issue where the `--no-config` option was not passed to child processes. This meant the child processes would still read the local config file and possibly cause unexpected behaviors.
- Fixed `info` command to eliminate “db (prior)” output if no backups or archives exist for a prior version of the cluster. ( *Fixed by Cynthia Shang. Reported by Stephen Frost.* )

Additional Notes

### Documentation Features:

- Document the relationship between the `archive-copy` and `archive-check` options. ( *Suggested by Markus Nullmeier.* )
- Improve `archive-copy` reference documentation.

## v1.26 Release Notes

Repository Encryption

Released November 21, 2017

### Bug Fixes:

- Fixed an issue that could cause copying large manifests to fail during restore. ( *Reported by Craig A. James.* )
- Fixed incorrect WAL offset for 32-bit architectures. ( *Fixed by Javier Wilson.* )
- Fixed an issue retrieving WAL for old database versions. After a stanza-upgrade it should still be possible to restore backups from the previous version and perform recovery with `archive-get` . However, `archive-get` only checked the most recent db version/id and failed. Also clean up some issues when the same db version/id appears multiple times in the history. ( *Fixed by Cynthia Shang. Reported by Clinton Adams.* )
- Fixed an issue with invalid backup groups being set correctly on restore. If the backup cannot map a group to a name it stores the group in the manifest as `false` then uses either the owner of `$PGDATA` to set the group during restore or failing that the group of the current user. This logic was not working correctly because the selected group was overwriting the user on restore leaving the group undefined and the user incorrectly set to the group. ( *Reported by Jeff McCormick.* )
- Fixed an issue passing parameters to remotes. When more than one db was specified the path, port, and socket path would for `db1` were passed no matter which db was actually being addressed. ( *Reported by uspen.* )

### Features:

- Repository encryption support. ( *Contributed by Cynthia Shang, David Steele.* )

### Improvements:

- Disable `gzip` filter when `--compress-level-network=0` . The filter was used with `compress` level set to 0 which added overhead without any benefit.
- Inflate performance improvement for `gzip` filter.

Additional Notes

### Documentation Features:

- Add template to improve initial information gathered for issue submissions. ( *Contributed by Cynthia Shang.* )

### Documentation Improvements:

- Clarify usage of the archive-timeout option and describe how it is distinct from the PostgreSQL archive\_timeout setting. ( *Contributed by Cynthia Shang. Suggested by Keith Fiske.* )

### **Test Suite Features:**

- Automated tests for 32-bit i386/i686 architecture.

### **v1.25 Release Notes**

#### S3 Performance Improvements

Released October 24, 2017

#### **Bug Fixes:**

- Fix custom settings for compress-level option being ignored. ( *Reported by Jens Wilke.* )
- Remove error when overlapping timelines are detected. Overlapping timelines are valid in many Point-in-Time-Recovery (PITR) scenarios. ( *Reported by blogh.* )
- Fix instances where database-id was not rendered as an integer in JSON info output. ( *Fixed by Cynthia Shang. Reported by Jason O'Donnell.* )

#### **Features:**

- Improve performance of list requests on S3. Any beginning literal portion of a filter expression is used to generate a search prefix which often helps keep the request small enough to avoid rate limiting. ( *Suggested by Mihail Shvein.* )

Additional Notes

### **Test Suite Features:**

- Add I/O performance tests.

### **v1.24 Release Notes**

#### New Backup Exclusions

Released September 28, 2017

#### **Bug Fixes:**

- Fixed an issue where warnings were being emitted in place of lower priority log messages during backup from standby initialization. ( *Reported by uspen.* )
- Fixed an issue where some db-\* options (e.g. db-port ) were not being passed to remotes. ( *Reported by uspen.* )

#### **Features:**

- Exclude contents of pg\_snapshots , pg\_serial , pg\_notify , and pg\_dynshmem from backup since they are rebuilt on startup.
- Exclude pg\_internal.init files from backup since they are rebuilt on startup.

#### **Improvements:**

- Open log file after async process is completely separated from the main process to prevent the main process from also logging to the file. ( *Suggested by Jens Wilke.* )

Additional Notes

### **Documentation Features:**

- Add passwordless SSH configuration.

### **Documentation Improvements:**

- Rename master to primary in documentation to align with PostgreSQL convention.

## v1.23 Release Notes

Multiple Standbys and PostgreSQL 10 Support

Released September 3, 2017

### Bug Fixes:

- Fixed an issue that could cause compression to abort on growing files. ( *Reported by Jesper St John, Aleksandr Rogozin.* )
- Fixed an issue with keep-alives not being sent to the remote from the local process. ( *Reported by William Cox.* )

### Features:

- Up to seven standbys can be configured for backup from standby. ( *Contributed by Cynthia Shang.* )
- PostgreSQL 10 support.
- Allow content-length (in addition to chunked encoding) when reading XML data to improve compatibility with third-party S3 gateways. ( *Suggested by Victor Gdalevich.* )

### Improvements:

- Increase HTTP timeout for S3.
- Add HTTP retries to harden against transient S3 network errors.

Additional Notes

### Documentation Bug Fixes:

- Fixed document generation to include section summaries on the Configuration page. ( *Fixed by Cynthia Shang.* )

## v1.22 Release Notes

Fixed S3 Retry

Released August 9, 2017

### Bug Fixes:

- Fixed authentication issue in S3 retry.

## v1.21 Release Notes

Improved Info Output and SSH Port Option

Released August 8, 2017

### Bug Fixes:

- The archive\_status directory is now recreated on restore to support PostgreSQL 8.3 which does not recreate it automatically like more recent versions do. ( *Reported by Stephen Frost.* )
- Fixed an issue that could cause the empty archive directory for an old PostgreSQL version to be left behind after a stanza-upgrade. ( *Fixed by Cynthia Shang.* )

### Features:

- Modified the info command (both text and JSON output) to display the archive ID and minimum/maximum WAL currently present in the archive for the current and prior, if any, database cluster version. ( *Contributed by Cynthia Shang.* )
- Added `-backup-ssh-port` and `-db-ssh-port` options to support non-default SSH ports. ( *Contributed by Cynthia Shang.* )

### Improvements:

- Retry when S3 returns an internal error (500).

Additional Notes

### Documentation Bug Fixes:

- Fix description of `--online` based on the command context.

### Documentation Features:

- Add creation of `/etc/pgbackrest.conf` to manual installation instructions.

### Documentation Improvements:

- Move repository options into a separate section in `command/command-line help`. ( *Suggested by Stephen Frost.* )

## v1.20 Release Notes

Critical 8.3/8.4 Bug Fix

Released June 27, 2017

**IMPORTANT NOTE** : PostgreSQL 8.3 and 8.4 installations utilizing tablespaces should upgrade immediately from any v1 release and run a full backup. A bug prevented tablespaces from being backed up on these versions only. PostgreSQL 9.0 is not affected.

### Bug Fixes:

- Fixed an issue that prevented tablespaces from being backed up on PostgreSQL 8.4 .
- Fixed missing flag in C library build that resulted in a mismatched binary on 32-bit systems. ( *Reported by Adrian Vondendriesch.* )

### Features:

- Add `s3-repo-ca-path` and `s3-repo-ca-file` options to accommodate systems where CAs are not automatically found by `IO::Socket::SSL` , i.e. RHEL7 , or to load custom CAs. ( *Suggested by Scott Frazer.* )

Additional Notes

### Test Suite Features:

- Add documentation builds to CI.

## v1.19 Release Notes

S3 Support

Released June 12, 2017

### Bug Fixes:

- Fixed the `info` command so the WAL archive min/max displayed is for the current database version. ( *Fixed by Cynthia Shang.* )
- Fixed the `backup` command so the `backup-standby` option is reset (and the backup proceeds on the primary) if the standby is not configured and/or reachable. ( *Fixed by Cynthia Shang.* )
- Fixed config warnings raised from a remote process causing errors in the master process. ( *Fixed by Cynthia Shang.* )

### Features:

- Amazon S3 repository support. ( *Reviewed by Cynthia Shang.* )

Additional Notes

### Documentation Bug Fixes:

- Changed invalid `max-archive-mb` option in configuration reference to `archive-queue-max` .
- Fixed missing `sudo` in installation section. ( *Fixed by Lætitia.* )

## v1.18 Release Notes

Stanza Upgrade, Refactoring, and Locking Improvements

Released April 12, 2017

### Bug Fixes:

- Fixed an issue where read-only operations that used local worker processes (i.e. restore) were creating write locks that could interfere with parallel archive-push. ( *Reported by Jens Wilke.* )

### Features:

- Added the stanza-upgrade command to provide a mechanism for upgrading a stanza after upgrading to a new major version of PostgreSQL. ( *Contributed by Cynthia Shang.* )
- Added validation of pgbackrest.conf to display warnings if options are not valid or are not in the correct section. ( *Contributed by Cynthia Shang.* )

### Improvements:

- Simplify locking scheme. Now, only the master process will hold write locks (for archive-push and backup commands) and not all local and remote worker processes as before.
- Do not set timestamps of files in the backup directories to match timestamps in the cluster directory. This was originally done to enable backup resume, but that process is now implemented with checksums.
- Improved error message when the restore command detects the presence of postmaster.pid. ( *Suggested by Yogesh Sharma.* )
- Renumber return codes between 25 and 125 to avoid PostgreSQL interpreting some as fatal signal exceptions. ( *Suggested by Yogesh Sharma.* )

## v1.17 Release Notes

Page Checksum Bug Fix

Released March 13, 2017

### Bug Fixes:

- Fixed an issue where newly initialized (but unused) pages would cause page checksum warnings. ( *Reported by Stephen Frost.* )

## v1.16 Release Notes

Page Checksum Improvements, CI, and Package Testing

Released March 2, 2017

### Bug Fixes:

- Fixed an issue where tables over 1GB would report page checksum warnings after the first segment. ( *Reported by Stephen Frost.* )
- Fixed an issue where databases created with a non-default tablespace would raise bogus warnings about pg\_filenode.map and pg\_internal.init not being page aligned. ( *Reported by blogh.* )

Additional Notes

### Test Suite Features:

- Continuous integration using travis-ci.
- Automated builds of Debian packages for all supported distributions.

## v1.15 Release Notes

Refactoring and Bug Fixes

Released February 13, 2017

### Bug Fixes:

- Fixed a regression introduced in v1.13 that could cause backups to fail if files were removed (e.g. tables dropped) while the manifest was being built. ( *Reported by Navid Golpayegani.* )

## v1.14 Release Notes

Refactoring and Bug Fixes

Released February 13, 2017

### Bug Fixes:

- Fixed an issue where an archive-push error would not be retried and would instead return errors to PostgreSQL indefinitely (unless the .error file was manually deleted). ( *Reported by Jens Wilke.* )
- Fixed a race condition in parallel archiving where creation of new paths generated an error when multiple processes attempted to do so at the same time. ( *Reported by Jens Wilke.* )

### Improvements:

- Improved performance of wal archive min/max provided by the info command. ( *Suggested by Jens Wilke.* )

Additional Notes

### Documentation Features:

- Updated async archiving documentation to more accurately describe how the new method works and how it differs from the old method. ( *Suggested by Jens Wilke.* )

## v1.13 Release Notes

Parallel Archiving, Stanza Create, Improved Info and Check

Released February 5, 2017

**IMPORTANT NOTE** : The new implementation of asynchronous archiving no longer copies WAL to a separate queue. If there is any WAL left over in the old queue after upgrading to 1.13 , it will be abandoned and **not** pushed to the repository.

To prevent this outcome, stop archiving by setting `archive_command = false` . Next, drain the async queue by running `pgbackrest -stanza=[stanza-name] archive-push` and wait for the process to complete. Check that the queue in `[spool-path]/archive/[stanza-name]/out` is empty. Finally, install 1.13 and restore the original `archive_command` .

**IMPORTANT NOTE** : The `stanza-create` command is not longer optional and must be executed before backup or archiving can be performed on a **new** stanza. Pre-existing stanzas do not require `stanza-create` to be executed.

### Bug Fixes:

- Fixed const assignment giving compiler warning in C library. ( *Fixed by Adrian Vondendriesch.* )
- Fixed a few directory syncs that were missed for the `-repo-sync` option.
- Fixed an issue where a missing user/group on restore could cause an “uninitialized value” error in `File->owner()` . ( *Reported by Leonardo GG Avellar.* )
- Fixed an issue where protocol mismatch errors did not output the expected value.
- Fixed a spurious `archive-get` log message that indicated an exit code of 1 was an abnormal termination.

### Features:

- Improved, multi-process implementation of asynchronous archiving.
- Improved `stanza-create` command so that it can repair broken repositories in most cases and is robust enough to be made mandatory. ( *Contributed by Cynthia Shang.* )
- Improved `check` command to run on a standby, though only basic checks are done because `pg_switch_xlog()` cannot be executed on a replica. ( *Contributed by Cynthia Shang.* )
- Added `archive` and `backup` WAL ranges to the `info` command.
- Added warning to update `pg_tablespace.splocation` when remapping tablespaces in PostgreSQL < 9.2. ( *Contributed by blogh.* )
- Remove remote lock requirements for the `archive-get` , `restore` , `info` , and `check` commands since they are read-only operations. ( *Suggested by Michael Vitale.* )

### Improvements:

- Log file banner is not output until the first log entry is written. ( *Suggested by Jens Wilke.* )

- Reduced the likelihood of torn pages causing a false positive in page checksums by filtering on start backup LSN.
- Remove Intel-specific optimization from C library build flags. ( *Contributed by Adrian Vondendriesch.* )
- Remove `-lock` option. This option was introduced before the lock directory could be located outside the repository and is now obsolete.
- Added `-log-timestamp` option to allow timestamps to be suppressed in logging. This is primarily used to avoid filters in the automated documentation.
- Return proper error code when unable to convert a relative path to an absolute path. ( *Suggested by Yogesh Sharma.* )

Additional Notes

#### Documentation Features:

- Added documentation to the User Guide for the `process-max` option. ( *Contributed by Cynthia Shang.* )

### v1.12 Release Notes

Page Checksums, Configuration, and Bug Fixes

Released December 12, 2016

**IMPORTANT NOTE** : In prior releases it was possible to specify options on the command-line that were invalid for the current command without getting an error. An error will now be generated for invalid options so it is important to carefully check command-line options in your environment to prevent disruption.

#### Bug Fixes:

- Fixed an issue where options that were invalid for the specified command could be provided on the command-line without generating an error. The options were ignored and did not cause any change in behavior, but it did lead to some confusion. Invalid options will now generate an error. ( *Reported by Nikhilchandra Kulkarni.* )
- Fixed an issue where internal symlinks were not being created for tablespaces in the repository. This issue was only apparent when trying to bring up clusters in-place manually using filesystem snapshots and did not affect normal backup and restore.
- Fixed an issue that prevented errors from being output to the console before the logging system was initialized, i.e. while parsing options. Error codes were still being returned accurately so this would not have made a process look like it succeeded when it did not. ( *Reported by Adrian Vondendriesch.* )
- Fixed an issue where the `db-port` option specified on the backup server would not be properly passed to the remote unless it was from the first configured database. ( *Reported by Michael Vitale.* )

#### Features:

- Added the `-checksum-page` option to allow `pgBackRest` to validate page checksums in data files when checksums are enabled on PostgreSQL  $\geq 9.3$ . Note that this functionality requires a C library which may not initially be available in OS packages. The option will automatically be enabled when the library is present and checksums are enabled on the cluster. ( *Suggested by Stephen Frost.* )
- Added the `-repo-link` option to allow internal symlinks to be suppressed when the repository is located on a filesystem that does not support symlinks. This does not affect any `pgBackRest` functionality, but the convenience link `latest` will not be created and neither will internal tablespace symlinks, which will affect the ability to bring up clusters in-place manually using filesystem snapshots.
- Added the `-repo-sync` option to allow directory syncs in the repository to be disabled for file systems that do not support them, e.g. NTFS.
- Added a predictable log entry to signal that a command has completed successfully. For example a backup ends successfully with: `INFO: backup command end: completed successfully .` ( *Suggested by Jens Wilke.* )

#### Improvements:

- For simplicity, the `pg_control` file is now copied with the rest of the files instead of by itself of at the end of the process. The backup command does not require this behavior and the restore copies to a temporary file which is renamed at the end of the restore.

Additional Notes

#### Documentation Bug Fixes:

- Fixed an issue that suppressed exceptions in PDF builds.
- Fixed regression in section links introduced in v1.10 .

#### Documentation Features:

- Added Retention to QuickStart section.

## v1.11 Release Notes

Bug Fix for Asynchronous Archiving Efficiency

Released November 17, 2016

### Bug Fixes:

- Fixed an issue where asynchronous archiving was transferring one file per execution instead of transferring files in batches. This regression was introduced in v1.09 and affected efficiency only, all WAL segments were correctly archived in asynchronous mode. ( *Reported by Stephen Frost.* )

## v1.10 Release Notes

Stanza Creation and Minor Bug Fixes

Released November 8, 2016

### Bug Fixes:

- Fixed an issue where a backup could error if no changes were made to a database between backups and only `pg_control` changed.
- Fixed an issue where tablespace paths with the same prefix would cause an invalid link error. ( *Reported by Nikhilchandra Kulkarni.* )

### Features:

- Added the `stanza-create` command to formalize creation of stanzas in the repository. ( *Contributed by Cynthia Shang.* )

### Improvements:

- Removed extraneous use lib directives from Perl modules. ( *Suggested by Devrim Gündüz.* )

## v1.09 Release Notes

9.6 Support, Configurability, and Bug Fixes

Released October 10, 2016

### Bug Fixes:

- Fixed the `check` command to prevent an error message from being logged if the backup directory does not exist. ( *Fixed by Cynthia Shang.* )
- Fixed error message to properly display the archive command when an invalid archive command is detected. ( *Reported by Jason O'Donnell.* )
- Fixed an issue where the async archiver would not be started if `archive-push` did not have enough space to queue a new WAL segment. This meant that the queue would never be cleared without manual intervention (such as calling `archive-push` directly). PostgreSQL now receives errors when there is not enough space to store new WAL segments but the async process will still be started so that space is eventually freed. ( *Reported by Jens Wilke.* )
- Fixed a remote timeout that occurred when a local process generated checksums (during resume or restore) but did not copy files, allowing the remote to go idle. ( *Reported by Jens Wilke.* )

### Features:

- Non-exclusive backups will automatically be used on PostgreSQL 9.6.
- Added the `cmd-ssh` option to allow the ssh client to be specified. ( *Suggested by Jens Wilke.* )
- Added the `log-level-stderr` option to control whether console log messages are sent to `stderr` or `stdout`. By default this is set to warn which represents a change in behavior from previous versions, even though it may be more intuitive. Setting `log-level-stderr=off` will preserve the old behavior. ( *Suggested by Sascha Biberhofer.* )
- Set `application_name` to “`pgBackRest [command]`” for database connections. ( *Suggested by Jens Wilke.* )
- Check that `archive_mode` is enabled when `archive-check` option enabled.

### Improvements:



- Clarified error message when unable to acquire pgBackRest advisory lock to make it clear that it is not a PostgreSQL backup lock. ( *Suggested by Jens Wilke.* )
- pgBackRest version number included in command start INFO log output.
- Process ID logged for local process start/stop INFO log output.

Additional Notes

#### Documentation Features:

- Added archive-timeout option documentation to the user guide. ( *Contributed by Cynthia Shang.* )

#### v1.08 Release Notes

Bug Fixes and Log Improvements

Released September 14, 2016

#### Bug Fixes:

- Fixed an issue where local processes were not disconnecting when complete and could later timeout. ( *Reported by Todd Vernick.* )
- Fixed an issue where the protocol layer could timeout while waiting for WAL segments to arrive in the archive. ( *Reported by Todd Vernick.* )

#### Improvements:

- Cache file log output until the file is created to create a more complete log.

#### v1.07 Release Notes

Thread to Process Conversion and Bug Fixes

Released September 7, 2016

#### Bug Fixes:

- Fixed an issue where tablespaces were copied from the primary during standby backup.
- Fixed the check command so backup info is checked remotely and not just locally. ( *Fixed by Cynthia Shang.* )
- Fixed an issue where retention-archive was not automatically being set when retention-archive-type=diff , resulting in a less aggressive than intended expiration of archive. ( *Fixed by Cynthia Shang.* )

#### Features:

- Converted Perl threads to processes to improve compatibility and performance.
- Exclude contents of \$PGDATA/pg\_replslot directory so that replication slots on the primary do not become part of the backup.
- The archive-start and archive-stop settings are now filled in backup.manifest even when archive-check=n . ( *Suggested by Jens Wilke.* )
- Additional warnings when archive retention settings may not have the intended effect or would allow indefinite retention. ( *Contributed by Cynthia Shang.* )
- Experimental support for non-exclusive backups in PostgreSQL 9.6 rc1. Changes to the control/catalog/WAL versions in subsequent release candidates may break compatibility but pgBackRest will be updated with each release to keep pace.

Additional Notes

#### Documentation Bug Fixes:

- Fixed minor documentation reproducibility issues related to binary paths.

#### Documentation Features:

- Documentation for archive retention. ( *Contributed by Cynthia Shang.* )

## v1.06 Release Notes

Backup from Standby and Bug Fixes

Released August 25, 2016

### Bug Fixes:

- Fixed an issue where a tablespace link that referenced another link would not produce an error, but instead skip the tablespace entirely. ( *Reported by Michael Vitale.* )
- Fixed an issue where options that should not allow multiple values could be specified multiple times in pgbackrest.conf without an error being raised. ( *Reported by Michael Vitale.* )
- Fixed an issue where the protocol-timeout option was not automatically increased when the db-timeout option was increased. ( *Reported by Todd Vernick.* )

### Features:

- Backup from a standby cluster. A connection to the primary cluster is still required to start/stop the backup and copy files that are not replicated, but the vast majority of files are copied from the standby in order to reduce load on the primary.
- More flexible configuration for databases. Master and standby can both be configured on the backup server and pgBackRest will automatically determine which is the primary. This means no configuration changes for backup are required after failing over from a primary to standby when a separate backup server is used.
- Exclude directories during backup that are cleaned, recreated, or zeroed by PostgreSQL at startup. These include `pgsql_tmp` and `pg_stat_tmp` . The `postgresql.auto.conf.tmp` file is now excluded in addition to files that were already excluded: `backup_label.old` , `postmaster.opts` , `postmaster.pid` , `recovery.conf` , `recovery.done` .
- Experimental support for non-exclusive backups in PostgreSQL 9.6 beta4. Changes to the control/catalog/WAL versions in subsequent betas may break compatibility but pgBackRest will be updated with each release to keep pace.

### Improvements:

- Improve error message for links that reference links in manifest build.
- Added hints to error message when relative paths are detected in `archive-push` or `archive-get` .
- Improve backup log messages to indicate which host the files are being copied from.

## v1.05 Release Notes

Bug Fix for Tablespace Link Checking

Released August 9, 2016

### Bug Fixes:

- Fixed an issue where tablespace paths that had `$PGDATA` as a substring would be identified as a subdirectories of `$PGDATA` even when they were not. Also hardened relative path checking a bit. ( *Reported by Chris Fort.* )

Additional Notes

### Documentation Features:

- Added documentation for scheduling backups with cron. ( *Contributed by Cynthia Shang.* )

### Documentation Improvements:

- Moved the backlog from the pgBackRest website to the GitHub repository wiki. ( *Contributed by Cynthia Shang.* )

## v1.04 Release Notes

Various Bug Fixes

Released July 30, 2016

### Bug Fixes:

- Fixed an issue where an extraneous remote was created causing threaded backup/restore to possibly timeout and/or throw a lock conflict. ( *Reported by Michael Vitale.* )
- Fixed an issue where db-path was not required for the check command so an assert was raised when it was missing rather than a polite error message. ( *Reported by Michael Vitale.* )
- Fixed check command to throw an error when database version/id does not match that of the archive. ( *Fixed by Cynthia Shang.* )
- Fixed an issue where a remote could try to start its own remote when the backup-host option was not present in pgbackrest.conf on the database server. ( *Reported by Lardière Sébastien.* )
- Fixed an issue where the contents of pg\_xlog were being backed up if the directory was symlinked. This didn't cause any issues during restore but was a waste of space.
- Fixed an invalid log() call in lock routines.

#### Features:

- Experimental support for non-exclusive backups in PostgreSQL 9.6 beta3. Changes to the control/catalog/WAL versions in subsequent betas may break compatibility but pgBackRest will be updated with each release to keep pace.

#### Improvements:

- Suppress banners on SSH protocol connections.
- Improved remote error messages to identify the host where the error was raised.
- All remote types now take locks. The exceptions date to when the test harness and pgBackRest were running in the same VM and no longer apply.

#### Additional Notes

#### Documentation Features:

- Added clarification on why the default for the backrest-user option is backrest . ( *Suggested by Michael Vitale.* )
- Updated information about package availability on supported platforms. ( *Suggested by Michael Vitale.* )

### v1.03 Release Notes

#### Check Command and Bug Fixes

Released July 2, 2016

#### Bug Fixes:

- Fixed an issue where keep-alives could be starved out by lots of small files during multi-threaded backup . They were also completely absent from single/multi-threaded backup resume and restore checksumming. ( *Reported by Janice Parkinson, Chris Barber.* )
- Fixed an issue where the expire command would refuse to run when explicitly called from the command line if the db-host option was set. This was not an issue when expire was run automatically after a backup ( *Reported by Chris Barber.* )
- Fixed an issue where validation was being running on archive\_command even when the archive-check option was disabled.

#### Features:

- Added check command to validate that pgBackRest is configured correctly for archiving and backups. ( *Contributed by Cynthia Shang.* )
- Added the protocol-timeout option. Previously protocol-timeout was set as db-timeout + 30 seconds.
- Failure to shutdown remotes at the end of the backup no longer throws an exception. Instead a warning is generated that recommends a higher protocol-timeout .
- Experimental support for non-exclusive backups in PostgreSQL 9.6 beta2. Changes to the control/catalog/WAL versions in subsequent betas may break compatibility but pgBackRest will be updated with each release to keep pace.

#### Improvements:

- Improved handling of users/groups captured during backup that do not exist on the restore host. Also explicitly handle the case where user/group is not mapped to a name.
- Option handling is now far more strict. Previously it was possible for a command to use an option that was not explicitly assigned to it. This was especially true for the backup-host and db-host options which are used to determine locality.

Additional Notes

### Documentation Improvements:

- Allow a static date to be used for documentation to generate reproducible builds. ( *Suggested by Adrian Vondendriesch.* )
- Added documentation for asynchronous archiving to the user guide. ( *Contributed by Cynthia Shang.* )
- Recommended install location for pgBackRest modules is now /usr/share/perl5 since /usr/lib/perl5 has been removed from the search path in newer versions of Perl.
- Added instructions for removing prior versions of pgBackRest .

### v1.02 Release Notes

Bug Fix for Perl 5.22

Released June 2, 2016

#### Bug Fixes:

- Fix usage of sprintf() due to new constraints in Perl 5.22. Parameters not referenced in the format string are no longer allowed. ( *Fixed by Adrian Vondendriesch.* )

Additional Notes

### Documentation Bug Fixes:

- Fixed syntax that was not compatible with Perl 5.2X. ( *Fixed by Christoph Berg, Adrian Vondendriesch.* )
- Fixed absolute paths that were used for the PDF logo. ( *Reported by Adrian Vondendriesch.* )

### Documentation Features:

- Release notes are now broken into sections so that bugs, features, and refactors are clearly delineated. An “Additional Notes” section has been added for changes to documentation and the test suite that do not affect the core code.
- Added man page generation. ( *Contributed by Adrian Vondendriesch, David Steele.* )
- The change log was the last piece of documentation to be rendered in Markdown only. Wrote a converter so the document can be output by the standard renderers. The change log will now be located on the website and has been renamed to “Releases” . ( *Contributed by Cynthia Shang.* )

### v1.01 Release Notes

Enhanced Info, Selective Restore, and 9.6 Support

Released May 17, 2016

#### Features:

- Enhanced text output of info command to include timestamps, sizes, and the reference list for all backups. ( *Contributed by Cynthia Shang.* )
- Allow selective restore of databases from a cluster backup. This feature can result in major space and time savings when only specific databases are restored. Unrestored databases will not be accessible but must be manually dropped before they will be removed from the shared catalogue. ( *Reviewed by Cynthia Shang, Greg Smith, Stephen Frost. Suggested by Stephen Frost.* )
- Experimental support for non-exclusive backups in PostgreSQL 9.6 beta1. Changes to the control/catalog/WAL versions in subsequent betas may break compatibility but pgBackRest will be updated with each release to keep pace. ( *Reviewed by Cynthia Shang.* )

### v1.00 Release Notes

New Repository Format and Configuration Scheme, Link Support

Released April 14, 2016

**IMPORTANT NOTE** : This flag day release breaks compatibility with older versions of pgBackRest . The manifest format, on-disk structure, configuration scheme, and the exe/path names have all changed. You must create a new repository to hold backups for this version of pgBackRest and keep your older repository for a time in case you need to do a restore. Restores from the prior repository will require the prior version of pgBackRest but because of name changes it is possible to have 1.00 and a prior version of pgBackRest installed at the same time. See the notes below for more detailed information on what has changed.

#### Features:

- Implemented a new configuration scheme which should be far simpler to use. See the User Guide and Configuration Reference for details but for a simple configuration all options can now be placed in the stanza section. Options that are shared between stanzas can be placed in the [global] section. More complex configurations can still make use of command sections though this should be a rare use case. ( *Suggested by Michael Renner.* )
- The repo-path option now always refers to the repository where backups and archive are stored, whether local or remote, so the repo-remote-path option has been removed. The new spool-path option can be used to define a location for queuing WAL segments when archiving asynchronously. A local repository is no longer required.
- The default configuration filename is now pgbackrest.conf instead of pg\_backrest.conf . This was done for consistency with other naming changes but also to prevent old config files from being loaded accidentally when migrating to 1.00 . ( *Suggested by Michael Renner, Stephen Frost.* )
- The default repository name was changed from /var/lib/backup to /var/lib/pgbackrest . ( *Suggested by Michael Renner, Stephen Frost.* )
- Lock files are now stored in /tmp/pgbackrest by default. These days /run/pgbackrest is the preferred location but that would require init scripts which are not part of this release. The lock-path option can be used to configure the lock directory.
- Log files are now stored in /var/log/pgbackrest by default and no longer have the date appended so they can be managed with logrotate . The log-path option can be used to configure the lock directory. ( *Suggested by Stephen Frost.* )
- Executable filename changed from pg\_backrest to pgbackrest . ( *Suggested by Michael Renner, Stephen Frost.* )
- All files and directories linked from PGDATA are now included in the backup. By default links will be restored directly into PGDATA as files or directories. The -link-all option can be used to restore all links to their original locations. The -link-map option can be used to remap a link to a new location.
- Removed -tablespace option and replaced with -tablespace-map-all option which should more clearly indicate its function.
- Added detail log level which will output more information than info without being as verbose as debug .

## Pre-Stable Releases

### v0.92 Release Notes

Command-line Repository Path Fix

Released April 6, 2016

#### Bug Fixes:

- Fixed an issue where the master process was passing -repo-remote-path instead of -repo-path to the remote and causing the lock files to be created in the default repository directory ( /var/lib/backup ), generally ending in failure. This was only an issue when -repo-remote-path was defined on the command line rather than in pg\_backrest.conf . ( *Reported by Jan Wieck.* )

### v0.91 Release Notes

Tablespace Bug Fix and Minor Enhancements

Released March 22, 2016

**IMPORTANT BUG FIX FOR TABLESPACES** : A change to the repository format was accidentally introduced in 0.90 which means the on-disk backup was no longer a valid PostgreSQL cluster when the backup contained tablespaces. This only affected users who directly copied the backups to restore PostgreSQL clusters rather than using the restore command. However, the fix breaks compatibility with older backups that contain tablespaces no matter how they are being restored ( pgBackRest will throw errors and refuse to restore). New full backups should be taken immediately after installing version 0.91 for any clusters that contain tablespaces. If older backups need to be restored then use a version of pgBackRest that matches the backup version.

#### Bug Fixes:

- Fixed repository incompatibility introduced in pgBackRest 0.90. ( *Reported by Evan Benoit.* )

#### Features:

- Copy global/pg\_control last during backups.
- Write .info and .manifest files to temp before moving them to their final locations and fsync'ing.
- Rename -no-start-stop option to -no-online .

Additional Notes

#### Test Suite Features:

- Static source analysis using Perl-Critic, currently passes on gentle.

## v0.90 Release Notes

### 9.5 Support, Various Enhancements, and Minor Bug Fixes

Released February 7, 2016

#### Bug Fixes:

- Fixed an issue where specifying `--no-archive-check` would throw a configuration error. ( *Reported by Jason O'Donnell.* )
- Fixed an issue where a temp WAL file left over after a well-timed system crash could cause the next archive-push to fail.
- The `retention-archive` option can now be safely set to less than backup retention ( `retention-full` or `retention-diff` ) without also specifying `archive-copy=n` . The WAL required to make the backups that fall outside of archive retention consistent will be preserved in the archive. However, in this case PITR will not be possible for the backups that fall outside of archive retention.

#### Features:

- When backing up and restoring tablespaces `pgBackRest` only operates on the subdirectory created for the version of PostgreSQL being run against. Since multiple versions can live in a tablespace (especially during a binary upgrade) this prevents too many files from being copied during a backup and other versions possibly being wiped out during a restore. This only applies to PostgreSQL  $\geq 9.0$  — prior versions of PostgreSQL could not share a tablespace directory.
- Generate an error when `archive-check=y` but `archive_command` does not execute `pg_backuprest` . ( *Contributed by Jason O'Donnell.* )
- Improved error message when `repo-path` or `repo-remote-path` does not exist.
- Added checks for `--delta` and `--force` restore options to ensure that the destination is a valid `$PGDATA` directory. `pgBackRest` will check for the presence of `PG_VERSION` or `backup.manifest` (left over from an aborted restore). If neither file is found then `--delta` and `--force` will be disabled but the restore will proceed unless there are files in the `$PGDATA` directory (or any tablespace directories) in which case the operation will be aborted.
- When restore `--set=latest` (the default) the actual backup restored will be output to the log.
- Support for PostgreSQL 9.5 partial WAL segments and `recovery_target_action` setting. The `archive_mode = 'always'` setting is not yet supported.
- Support for `recovery_target = 'immediate'` recovery setting introduced in PostgreSQL 9.4.
- The following tablespace checks have been added: paths or files in `pg_tblspc`, relative links in `pg_tblspc`, tablespaces in `$PGDATA`. All three will generate errors.

## v0.89 Release Notes

### Timeout Bug Fix and Restore Read-Only Repositories

Released December 24, 2015

#### Bug Fixes:

- Fixed an issue where longer-running backups/restores would timeout when remote and threaded. Keepalives are now used to make sure the remote for the main process does not timeout while the thread remotes do all the work. The error message for timeouts was also improved to make debugging easier. ( *Reported by Stephen Frost.* )

#### Features:

- Allow restores to be performed on a read-only repository by using `--no-lock` and `--log-level-file=off` . The `--no-lock` option can only be used with restores.

## v0.88 Release Notes

### Documentation and Minor Bug Fixes

Released November 22, 2015

#### Bug Fixes:

- Fixed an issue where the start / stop commands required the `--config` option. ( *Reported by Dmitry Didovicher.* )
- Fixed an issue where log files were being overwritten instead of appended. ( *Reported by Stephen Frost, Dmitry Didovicher.* )
- Fixed an issue where `backup-user` was not optional.

#### Features:

- Symlinks are no longer created in backup directories in the repository. These symlinks could point virtually anywhere and potentially be dangerous. Symlinks are still recreated during a restore. ( *Suggested by Stephen Frost.* )
- Added better messaging for backup expiration. Full and differential backup expirations are logged on a single line along with a list of all dependent backups expired.
- Archive retention is automatically set to full backup retention if not explicitly configured.

Additional Notes

### Documentation Features:

- Added documentation in the user guide for delta restores, expiration, dedicated backup hosts, starting and stopping pgBackRest, and replication.

### v0.87 Release Notes

Website and User Guide

Released October 28, 2015

### Features:

- The backup\_label.old and recovery.done files are now excluded from backups.

Additional Notes

### Documentation Features:

- Added a new user guide that covers pgBackRest basics and some advanced topics including PITR. Much more to come, but it's a start. ( *Contributed by David Steele, Stephen Frost. Reviewed by Michael Renner, Cynthia Shang, Eric Radman, Dmitry Didovicher.* )

### v0.85 Release Notes

Start/Stop Commands and Minor Bug Fixes

Released October 8, 2015

### Bug Fixes:

- Fixed an issue where an error could be returned after a backup or restore completely successfully.
- Fixed an issue where a resume would fail if temp files were left in the root backup directory when the backup failed. This scenario was likely if the backup process got terminated during the copy phase.

### Features:

- Added stop and start commands to prevent pgBackRest processes from running on a system where PostgreSQL is shutdown or the system needs to be quiesced for some other reason.
- Experimental support for PostgreSQL 9.5 beta1. This may break when the control version or WAL magic changes in future versions but will be updated in each pgBackRest release to keep pace. All regression tests pass except for `-target-resume` tests (this functionality has changed in 9.5) and there is no testing yet for `.partial` WAL segments.

### v0.82 Release Notes

Refactoring, Command-line Help, and Minor Bug Fixes

Released September 14, 2015

### Bug Fixes:

- Fixed an issue where resumed compressed backups were not preserving existing files.
- Fixed an issue where resume and incr/diff would not ensure that the prior backup had the same compression and hardlink settings.
- Fixed an issue where a cold backup using `-no-start-stop` could be started on a running PostgreSQL cluster without `-force` specified.
- Fixed an issue where a thread could be started even when none were requested.

- Fixed an issue where the pgBackRest version number was not being updated in backup.info and archive.info after an upgrade/downgrade.
- Fixed an issue where the info command was throwing an exception when the repository contained no stanzas. ( *Reported by Stephen Frost.* )
- Fixed an issue where the PostgreSQL pg\_stop\_backup() NOTICES were being output to stderr . ( *Reported by Stephen Frost.* )

### Features:

- Experimental support for PostgreSQL 9.5 alpha2. This may break when the control version or WAL magic changes in future versions but will be updated in each pgBackRest release to keep pace. All regression tests pass except for `-target-resume` tests (this functionality has changed in 9.5) and there is no testing yet for `.partial` WAL segments.

### Improvements:

- Renamed recovery-setting option and section to `recovery-option` to be more consistent with pgBackRest naming conventions.
- Added dynamic module loading to speed up commands, especially asynchronous archiving.

### Additional Notes

### Documentation Features:

- Command-line help is now extracted from the same XML source that is used for the other documentation and includes much more detail.

### v0.80 Release Notes

#### DBI Support, Stability, and Convenience Features

Released August 9, 2015

### Bug Fixes:

- Fixed an issue that caused the formatted timestamp for both the oldest and newest backups to be reported as the current time by the info command. Only text output was affected – json output reported the correct epoch values. ( *Reported by Michael Renner.* )
- Fixed protocol issue that was preventing ssh errors (especially on connection) from being logged.

### Features:

- The repository is now created and updated with consistent directory and file modes. By default `umask` is set to `0000` but this can be disabled with the `neutral-umask` setting. ( *Suggested by Cynthia Shang.* )
- Added the `stop-auto` option to allow failed backups to automatically be stopped when a new backup starts.
- Added the `db-timeout` option to limit the amount of time pgBackRest will wait for `pg_start_backup()` and `pg_stop_backup()` to return.
- Remove `pg_control` file at the beginning of the restore and copy it back at the very end. This prevents the possibility that a partial restore can be started by PostgreSQL .
- Added checks to be sure the `db-path` setting is consistent with `db-port` by comparing the `data_directory` as reported by the cluster against the `db-path` setting and the version as reported by the cluster against the value read from `pg_control` . The `db-socket-path` setting is checked to be sure it is an absolute path.
- Experimental support for PostgreSQL 9.5 alpha1. This may break when the control version or WAL magic changes in future versions but will be updated in each pgBackRest release to keep pace. All regression tests pass except for `-target-resume` tests (this functionality has changed in 9.5) and there is no testing yet for `.partial` WAL segments.

### Improvements:

- Now using Perl DBI and `DBD::Pg` for connections to PostgreSQL rather than `psql` . The `cmd-psql` and `cmd-psql-option` settings have been removed and replaced with `db-port` and `db-socket-path` . Follow the instructions in the Installation Guide to install `DBD::Pg` on your operating system.

### Additional Notes

### Test Suite Features:

- Added vagrant test configurations for Ubuntu 14.04 and CentOS 7.



## v0.78 Release Notes

Remove CPAN Dependencies, Stability Improvements

Released July 13, 2015

### Improvements:

- Removed dependency on CPAN packages for multi-threaded operation. While it might not be a bad idea to update the threads and Thread::Queue packages, it is no longer necessary.
- Modified wait backoff to use a Fibonacci rather than geometric sequence. This will make wait time grow less aggressively while still giving reasonable values.

Additional Notes

### Test Suite Features:

- Added vagrant test configurations for Ubuntu 12.04 and CentOS 6.

## v0.77 Release Notes

CentOS/RHEL 6 Support and Protocol Improvements

Released June 30, 2015

### Features:

- Added file and directory syncs to the File object for additional safety during backup/restore and archiving. ( *Suggested by Andres Freund.* )
- Added support for Perl 5.10.1 and OpenSSH 5.3 which are default for CentOS/RHEL 6. ( *Suggested by Eric Radman.* )
- Improved error message when backup is run without archive\_command set and without --no-archive-check specified. ( *Suggested by Eric Radman.* )

## v0.75 Release Notes

New Repository Format, Info Command and Experimental 9.5 Support

Released June 14, 2015

**IMPORTANT NOTE** : This flag day release breaks compatibility with older versions of pgBackRest . The manifest format, on-disk structure, and the binary names have all changed. You must create a new repository to hold backups for this version of pgBackRest and keep your older repository for a time in case you need to do a restore. The pg\_backrest.conf file has not changed but you'll need to change any references to pg\_backrest.pl in cron (or elsewhere) to pg\_backrest (without the .pl extension).

### Features:

- Added the info command.
- Logging now uses unbuffered output. This should make log files that are being written by multiple threads less chaotic. ( *Suggested by Michael Renner.* )
- Experimental support for PostgreSQL 9.5. This may break when the control version or WAL magic changes but will be updated in each release.

### Improvements:

- More efficient file ordering for backup . Files are copied in descending size order so a single thread does not end up copying a large file at the end. This had already been implemented for restore .

## v0.70 Release Notes

Stability Improvements for Archiving, Improved Logging and Help

Released June 1, 2015

### Bug Fixes:

- Fixed an issue where archive-copy would fail on an incr/diff backup when hardlink=n . In this case the pg\_xlog path does not already exist and must be created. ( *Reported by Michael Renner.* )
- Fixed an issue in async archiving where archive-push was not properly returning 0 when archive-max-mb was reached and moved the async check after transfer to avoid having to remove the stop file twice. Also added unit tests for this case and improved error messages to make it clearer to the user what went wrong. ( *Reported by Michael Renner.* )
- Fixed a locking issue that could allow multiple operations of the same type against a single stanza. This appeared to be benign in terms of data integrity but caused spurious errors while archiving and could lead to errors in backup/restore. ( *Reported by Michael Renner.* )

#### Features:

- Allow duplicate WAL segments to be archived when the checksum matches. This is necessary for some recovery scenarios.
- Allow comments/disabling in pg\_backrest.conf using the # character. Only # characters in the first character of the line are honored. ( *Suggested by Michael Renner.* )
- Better logging before pg\_start\_backup() to make it clear when the backup is waiting on a checkpoint. ( *Suggested by Michael Renner.* )
- Various command behavior and logging fixes. ( *Reviewed by Michael Renner. Suggested by Michael Renner.* )

#### Improvements:

- Replaced JSON module with JSON::PP which ships with core Perl.

#### Additional Notes

#### Documentation Bug Fixes:

- Various help fixes. ( *Reviewed by Michael Renner. Reported by Michael Renner.* )

#### v0.65 Release Notes

##### Improved Resume and Restore Logging, Compact Restores

Released May 11, 2015

#### Bug Fixes:

- Fixed an issue where an absolute path was not written into recovery.conf when the restore was run with a relative path.

#### Features:

- Better resume support. Resumed files are checked to be sure they have not been modified and the manifest is saved more often to preserve checksums as the backup progresses. More unit tests to verify each resume case.
- Resume is now optional. Use the resume setting or -no-resume from the command line to disable.
- More info messages during restore. Previously, most of the restore messages were debug level so not a lot was output in the log.
- Added tablespace setting to allow tablespaces to be restored into the pg\_tblspc path. This produces compact restores that are convenient for development, staging, etc. Currently these restores cannot be backed up as pgBackRest expects only links in the pg\_tblspc path.

#### v0.61 Release Notes

##### Bug Fix for Uncompressed Remote Destination

Released April 21, 2015

#### Bug Fixes:

- Fixed a buffering error that could occur on large, highly-compressible files when copying to an uncompressed remote destination. The error was detected in the decompression code and resulted in a failed backup rather than corruption so it should not affect successful backups made with previous versions.

## v0.60 Release Notes

Better Version Support and WAL Improvements

Released April 19, 2015

### Bug Fixes:

- Pushing duplicate WAL now generates an error. This worked before only if checksums were disabled.

### Features:

- Database System IDs are used to make sure that all WAL in an archive matches up. This should help prevent misconfigurations that send WAL from multiple clusters to the same archive.

Additional Notes

### Test Suite Features:

- Regression tests working back to PostgreSQL 8.3.

## v0.50 Release Notes

Restore and Much More

Released March 25, 2015

### Bug Fixes:

- Fixed broken checksums and now they work with normal and resumed backups. Finally realized that checksums and checksum deltas should be functionally separated and this simplified a number of things. Issue #28 has been created for checksum deltas.
- Fixed an issue where a backup could be resumed from an aborted backup that didn't have the same type and prior backup.

### Features:

- Added restore functionality.
- All options can now be set on the command-line making `pg_backrest.conf` optional.
- De/compression is now performed without threads and checksum/size is calculated in stream. That means file checksums are no longer optional.
- Added option `-no-start-stop` to allow backups when Postgres is shut down. If `postmaster.pid` is present then `-force` is required to make the backup run (though if Postgres is running an inconsistent backup will likely be created). This option was added primarily for the purpose of unit testing, but there may be applications in the real world as well.
- Checksum for `backup.manifest` to detect a corrupted/modified manifest.
- Link latest always points to the last backup. This has been added for convenience and to make restores simpler.

Additional Notes

### Test Suite Features:

- More comprehensive unit tests in all areas.

## v0.30 Release Notes

Core Restructuring and Unit Tests

Released October 5, 2014

Additional Notes

### Documentation Features:

- Added much needed documentation

### Test Suite Features:

- Fairly comprehensive unit tests for all the basic operations. More work to be done here for sure, but then there is always more work to be done on unit tests.

## v0.19 Release Notes

Improved Error Reporting/Handling

Released May 13, 2014

### Bug Fixes:

- Found and squashed a nasty bug where `file_copy()` was defaulted to ignore errors. There was also an issue in `file_exists()` that was causing the test to fail when the file actually did exist. Together they could have resulted in a corrupt backup with no errors, though it is very unlikely.

## v0.18 Release Notes

Return Soft Error When Archive Missing

Released April 13, 2014

### Bug Fixes:

- The `archive-get` command now returns a 1 when the archive file is missing to differentiate from hard errors (ssh connection failure, file copy error, etc.) This lets PostgreSQL know that the archive stream has terminated normally. However, this does not take into account possible holes in the archive stream. ( *Reported by Stephen Frost.* )

## v0.17 Release Notes

Warn When Archive Directories Cannot Be Deleted

Released April 3, 2014

### Bug Fixes:

- If an archive directory which should be empty could not be deleted backrest was throwing an error. There's a good fix for that coming, but for the time being it has been changed to a warning so processing can continue. This was impacting backups as sometimes the final archive file would not get pushed if the first archive file had been in a different directory (plus some bad luck).

## v0.16 Release Notes

RequestTTY=yes for SSH Sessions

Released April 1, 2014

### Bug Fixes:

- Added `RequestTTY=yes` to ssh sessions. Hoping this will prevent random lockups.

## v0.15 Release Notes

Added `archive-get`

Released March 29, 2014

### Features:

- Added `archive-get` functionality to aid in restores.
- Added option to force a checkpoint when starting the backup, `start-fast=y` .

## v0.11 Release Notes

Minor Fixes

Released March 26, 2014

### Bug Fixes:

- Removed `master_stderr_discard` option on database SSH connections. There have been occasional lockups and they could be related to issues originally seen in the file code. ( *Reported by Stephen Frost.* )
- Changed lock file conflicts on backup and expire commands to `ERROR` . They were set to `DEBUG` due to a copy-and-paste from the archive locks.

## v0.10 Release Notes

Backup and Archiving are Functional

Released March 5, 2014

### Features:

- No restore functionality, but the backup directories are consistent PostgreSQL data directories. You'll need to either uncompress the files or turn off compression in the backup. Uncompressed backups on a ZFS (or similar) filesystem are a good option because backups can be restored locally via a snapshot to create logical backups or do spot data recovery.
- Archiving is single-threaded. This has not posed an issue on our multi-terabyte databases with heavy write volume. Recommend a large WAL volume or to use the async option with a large volume nearby.
- Backups are multi-threaded, but the Net::OpenSSH library does not appear to be 100% thread-safe so it will very occasionally lock up on a thread. There is an overall process timeout that resolves this issue by killing the process. Yes, very ugly.
- Checksums are lost on any resumed backup. Only the final backup will record checksum on multiple resumes. Checksums from previous backups are correctly recorded and a full backup will reset everything.
- The backup.manifest is being written as Storable because Config::IniFile does not seem to handle large files well. Would definitely like to save these as human-readable text.

Additional Notes

### Documentation Features:

- Absolutely no documentation (outside the code). Well, excepting these release notes.