



# **Trusted Postgres Architect**

***Release 23.41.0***

**EnterpriseDB Corporation**

**Jan 05, 2026**



|          |                                                  |          |
|----------|--------------------------------------------------|----------|
| <b>1</b> | <b>Introduction</b>                              | <b>3</b> |
| 1.1      | Introduction                                     | 3        |
| 1.1.1    | What can TPA do?                                 | 3        |
| 1.1.2    | How do I use it?                                 | 4        |
| 1.1.3    | It's just Postgres                               | 6        |
| 1.1.4    | Versioning in TPA                                | 6        |
| 1.1.5    | Getting started                                  | 6        |
| 1.2      | Release notes                                    | 6        |
| 1.2.1    | Trusted Postgres Architect 23.41.0 release notes | 7        |
| 1.2.2    | Trusted Postgres Architect 23.40.1 release notes | 12       |
| 1.2.3    | Trusted Postgres Architect 23.40.0 release notes | 13       |
| 1.2.4    | Trusted Postgres Architect 23.39.0 release notes | 13       |
| 1.2.5    | Trusted Postgres Architect 23.38.1 release notes | 16       |
| 1.2.6    | Trusted Postgres Architect 23.38.0 release notes | 17       |
| 1.2.7    | Trusted Postgres Architect 23.37.0 release notes | 20       |
| 1.2.8    | Trusted Postgres Architect 23.36.0 release notes | 22       |
| 1.2.9    | Trusted Postgres Architect 23.35.0 release notes | 25       |
| 1.2.10   | Trusted Postgres Architect 23.34.1 release notes | 28       |
| 1.2.11   | Trusted Postgres Architect 23.34 release notes   | 29       |
| 1.2.12   | Trusted Postgres Architect 23.33 release notes   | 30       |
| 1.2.13   | Trusted Postgres Architect 23.32 release notes   | 31       |
| 1.2.14   | Trusted Postgres Architect 23.31 release notes   | 32       |
| 1.2.15   | Trusted Postgres Architect 23.30 release notes   | 32       |
| 1.2.16   | Trusted Postgres Architect 23.29 release notes   | 33       |
| 1.2.17   | Trusted Postgres Architect 23.28 release notes   | 34       |
| 1.2.18   | Trusted Postgres Architect 23.27 release notes   | 35       |
| 1.2.19   | Trusted Postgres Architect 23.26 release notes   | 35       |
| 1.2.20   | Trusted Postgres Architect 23.25 release notes   | 36       |
| 1.2.21   | Trusted Postgres Architect 23.24 release notes   | 36       |
| 1.2.22   | Trusted Postgres Architect 23.23 release notes   | 37       |
| 1.2.23   | Trusted Postgres Architect 23.22 release notes   | 37       |
| 1.2.24   | Trusted Postgres Architect 23.21 release notes   | 38       |
| 1.2.25   | Trusted Postgres Architect 23.20 release notes   | 38       |
| 1.2.26   | Trusted Postgres Architect 23.19 release notes   | 39       |
| 1.2.27   | Trusted Postgres Architect 23.18 release notes   | 39       |
| 1.2.28   | Trusted Postgres Architect 23.17 release notes   | 40       |

|          |                                                                |           |
|----------|----------------------------------------------------------------|-----------|
| 1.2.29   | Trusted Postgres Architect 23.16 release notes                 | 40        |
| 1.2.30   | Trusted Postgres Architect 23.15 release notes                 | 41        |
| 1.2.31   | Trusted Postgres Architect 23.14 release notes                 | 41        |
| 1.2.32   | Trusted Postgres Architect 23.13 release notes                 | 41        |
| 1.2.33   | Trusted Postgres Architect 23.12 release notes                 | 41        |
| 1.2.34   | Trusted Postgres Architect 23.1 to 23.11 release notes         | 42        |
| 1.3      | TPA installation                                               | 48        |
| 1.3.1    | Quickstart                                                     | 49        |
| 1.3.2    | Where to install TPA                                           | 49        |
| 1.3.3    | Installing TPA packages                                        | 49        |
| 1.3.4    | Setting up the TPA Python environment                          | 50        |
| 1.3.5    | Installing TPA without internet or network access (air-gapped) | 50        |
| 1.3.6    | Verifying your TPA installation                                | 51        |
| 1.3.7    | Upgrading TPA                                                  | 51        |
| 1.3.8    | Ansible versions                                               | 51        |
| 1.4      | Open source TPA                                                | 51        |
| 1.4.1    | What is Trusted Postgres Architect (TPA)?                      | 51        |
| 1.4.2    | Next Steps                                                     | 51        |
| 1.4.3    | TPA Open Source FAQs                                           | 52        |
| 1.5      | Installing TPA from source                                     | 52        |
| 1.5.1    | Quickstart                                                     | 52        |
| 1.5.2    | Clone and setup                                                | 52        |
| 1.5.3    | Dependencies                                                   | 53        |
| <b>2</b> | <b>Creating a cluster</b>                                      | <b>55</b> |
| 2.1      | A First Cluster Deployment                                     | 55        |
| 2.1.1    | Installing TPA                                                 | 55        |
| 2.1.2    | Installing Docker                                              | 55        |
| 2.1.3    | Creating a configuration with TPA                              | 56        |
| 2.2      | Cluster configuration                                          | 58        |
| 2.2.1    | Quickstart                                                     | 59        |
| 2.2.2    | Configuration options                                          | 59        |
| 2.2.3    | Software selection                                             | 61        |
| 2.2.4    | Overrides                                                      | 63        |
| 2.2.5    | Ansible Tower                                                  | 64        |
| 2.2.6    | PGAI agent                                                     | 64        |
| 2.2.7    | Git repository                                                 | 64        |
| 2.2.8    | Keyring backend for vault password                             | 64        |
| 2.2.9    | Security standards compliance                                  | 64        |
| 2.2.10   | Examples                                                       | 65        |
| 2.3      | tpaexec deploy                                                 | 67        |
| 2.3.1    | Prerequisites                                                  | 67        |
| 2.3.2    | Quickstart                                                     | 67        |
| 2.3.3    | Selective deployment                                           | 68        |
| 2.3.4    | deploy.yml                                                     | 68        |
| 2.4      | tpaexec provision                                              | 68        |
| 2.4.1    | Prerequisites                                                  | 68        |
| 2.4.2    | Quickstart                                                     | 69        |
| 2.4.3    | Options                                                        | 69        |
| 2.4.4    | Accessing the instances                                        | 69        |
| 2.4.5    | Generated files                                                | 70        |
| 2.5      | tpaexec test                                                   | 71        |
| 2.5.1    | Quickstart                                                     | 72        |

|          |                             |           |
|----------|-----------------------------|-----------|
| <b>3</b> | <b>Architectures</b>        | <b>73</b> |
| 3.1      | PGD-Always-ON               | 73        |
| 3.1.1    | Cluster configuration       | 73        |
| 3.2      | PGD Lightweight             | 75        |
| 3.2.1    | Cluster configuration       | 75        |
| 3.3      | BDR-Always-ON               | 76        |
| 3.3.1    | Cluster configuration       | 77        |
| 3.4      | M1                          | 78        |
| 3.4.1    | Failover management         | 78        |
| 3.4.2    | Cluster configuration       | 79        |
| 3.5      | PGD-S                       | 80        |
| 3.5.1    | Cluster configuration       | 80        |
| 3.6      | PGD-X                       | 81        |
| 3.6.1    | Cluster configuration       | 82        |
| <b>4</b> | <b>Platforms</b>            | <b>85</b> |
| 4.1      | aws                         | 85        |
| 4.1.1    | API access setup            | 85        |
| 4.1.2    | Introduction                | 87        |
| 4.1.3    | Configuration               | 87        |
| 4.2      | bare(-metal servers)        | 91        |
| 4.2.1    | SSH access                  | 91        |
| 4.2.2    | Distribution support        | 92        |
| 4.2.3    | IP addresses                | 92        |
| 4.2.4    | Starting afresh             | 92        |
| 4.3      | Docker                      | 93        |
| 4.3.1    | Synopsis                    | 93        |
| 4.3.2    | Operating system selection  | 93        |
| 4.3.3    | Installing Docker           | 93        |
| 4.3.4    | Docker container management | 95        |
| <b>5</b> | <b>Customizations</b>       | <b>97</b> |
| 5.1      | Cluster configuration       | 97        |
| 5.1.1    | config.yml                  | 97        |
| 5.1.2    | Variables                   | 98        |
| 5.1.3    | Cluster variables           | 98        |
| 5.1.4    | Instance variables          | 99        |
| 5.1.5    | instance_defaults           | 100       |
| 5.1.6    | Locations                   | 100       |
| 5.2      | Instance configuration      | 101       |
| 5.2.1    | System-level configuration  | 101       |
| 5.2.2    | Postgres                    | 102       |
| 5.2.3    | Other components            | 103       |
| 5.2.4    | After starting Postgres     | 103       |
| 5.3      | Building from source        | 104       |
| 5.3.1    | Quickstart                  | 104       |
| 5.3.2    | Configuration               | 105       |
| 5.3.3    | Local source directories    | 105       |
| 5.3.4    | Rebuilding                  | 106       |
| 5.4      | TPA hooks                   | 106       |
| 5.4.1    | Summary                     | 106       |
| 5.4.2    | General-purpose hooks       | 107       |
| 5.4.3    | PGD hooks                   | 108       |
| 5.4.4    | Other hooks                 | 109       |

|          |                                                   |            |
|----------|---------------------------------------------------|------------|
| 5.4.5    | New hooks                                         | 109        |
| <b>6</b> | <b>Cluster management</b>                         | <b>111</b> |
| 6.1      | Upgrading your cluster                            | 111        |
| 6.1.1    | Component selection                               | 112        |
| 6.1.2    | Package version selection                         | 113        |
| 6.1.3    | Configuration                                     | 114        |
| 6.1.4    | Upgrading from BDR-Always-ON to PGD-Always-ON     | 114        |
| 6.1.5    | Upgrading from PGD-Always-ON to PGD-X             | 115        |
| 6.1.6    | PGD-S or PGD-X                                    | 117        |
| 6.1.7    | PGD-Always-ON                                     | 118        |
| 6.1.8    | BDR-Always-ON                                     | 118        |
| 6.1.9    | M1                                                | 119        |
| 6.1.10   | Controlling the upgrade process                   | 119        |
| 6.1.11   | Upgrading a Subset of Nodes                       | 119        |
| 6.2      | tpaexec switchover                                | 120        |
| 6.2.1    | Example                                           | 120        |
| 6.2.2    | Architecture options                              | 120        |
| 6.2.3    | Repmgr redirect pgbouncer                         | 120        |
| 6.3      | BDR/HAProxy server pool management                | 120        |
| 6.3.1    | Example                                           | 121        |
| 6.4      | tpaexec rehydrate                                 | 121        |
| 6.4.1    | Prerequisites                                     | 121        |
| 6.4.2    | Example                                           | 122        |
| 6.4.3    | Rehydrate in phases                               | 123        |
| 6.4.4    | Appendix                                          | 123        |
| <b>7</b> | <b>Miscellaneous</b>                              | <b>125</b> |
| 7.1      | TPA and Ansible Tower/Ansible Automation Platform | 125        |
| 7.1.1    | AAP initial setup                                 | 125        |
| 7.1.2    | Setting up a cluster                              | 126        |
| 7.1.3    | Use one project for multiple inventory            | 127        |
| 7.1.4    | Update TPA on AAP                                 | 128        |
| 7.1.5    | Build an EE for TPA                               | 128        |
| 7.2      | TPA, Ansible, and sudo                            | 129        |
| 7.2.1    | Ansible sudo invocations                          | 130        |
| 7.2.2    | Recommendations                                   | 130        |
| 7.2.3    | SSH and sudo passwords                            | 130        |
| 7.2.4    | sudo options                                      | 131        |
| 7.2.5    | Logging                                           | 131        |
| 7.2.6    | Local privileges                                  | 131        |
| 7.3      | TPA - PuTTY Configuration guide                   | 131        |
| 7.3.1    | Key conversion                                    | 131        |
| 7.3.2    | <b>Configure PuTTY</b>                            | 132        |
| 7.4      | Troubleshooting                                   | 132        |
| 7.4.1    | Recreate python virtual environment               | 132        |
| 7.4.2    | Strange AWS errors regarding credentials          | 133        |
| 7.4.3    | Logging                                           | 133        |
| 7.4.4    | Cluster test                                      | 133        |
| 7.4.5    | TPA server test                                   | 134        |
| 7.4.6    | Including or excluding specific tasks             | 134        |
| 7.5      | Reference                                         | 134        |
| 7.5.1    | Installation                                      | 134        |
| 7.5.2    | Administration                                    | 138        |

|       |                                                               |     |
|-------|---------------------------------------------------------------|-----|
| 7.5.3 | Services . . . . .                                            | 141 |
| 7.5.4 | Configuration . . . . .                                       | 172 |
| 7.5.5 | Postgres configuration . . . . .                              | 190 |
| 7.5.6 | Commands . . . . .                                            | 194 |
| 7.6   | Selective task execution . . . . .                            | 219 |
| 7.6.1 | Using task selectors . . . . .                                | 219 |
| 7.6.2 | Examples . . . . .                                            | 220 |
| 7.6.3 | Supported selectors for <code>tpaexec deploy</code> . . . . . | 220 |
| 7.6.4 | Supported selectors for <code>tpaexec test</code> . . . . .   | 224 |



© Copyright EnterpriseDB UK Limited 2015-2025 - All rights reserved.



## 1.1 Introduction

TPA is an orchestration tool that uses Ansible to deploy Postgres clusters according to EDB's recommendations.

TPA embodies the best practices followed by EDB, informed by many years of hard-earned experience with deploying and supporting Postgres. These recommendations are as applicable to quick testbed setups as to production environments.

### 1.1.1 What can TPA do?

TPA is built around a declarative configuration mechanism that you can use to describe a Postgres cluster, from its topology right down to the smallest details of its configuration.

Start by running `tpaexec configure` to generate an initial cluster configuration based on a few high-level choices (e.g., which version of Postgres to install). The default configuration is ready to use as-is, but you can edit it to suit your needs (the generated configuration is just a text file, `config.yml`).

Using this configuration, TPA can:

1. Provision servers (e.g., AWS EC2 instances or Docker containers) and any other resources needed to host the cluster (or you can deploy to existing servers or VMs just by specifying connection details).
2. Configure the operating system (tweak kernel settings, create users and SSH keys, install packages, define systemd services, set up log rotation, and so on).
3. Install and configure Postgres and associated components (e.g., PGD, Barman, pgbouncer, repmgr, and various Postgres extensions).
4. Run automated tests on the cluster after deployment.
5. Deploy future changes to your configuration (e.g., changing Postgres settings, installing and upgrading packages, adding new servers, and so on).

### 1.1.2 How do I use it?

To use TPA, you need to install it and run the `tpaexec setup` command. Follow the *TPA installation* for your platform.

TPA operates in four distinct stages to bring up a Postgres cluster:

- Generate a cluster *Cluster configuration*
- *Provisioning the deployment* servers (VMs, containers) to host the cluster
- *A First Cluster Deployment* software to the provisioned instances
- *Testing* the deployed cluster

```
# 1. Configuration: decide what kind of cluster you want

tpaexec configure clustername --architecture M1 --platform aws \
    --postgresql 14 \
    --failover-manager repmgr

# 2. Provisioning: create the servers needed to host the cluster

tpaexec provision clustername

# 3. Deployment: install and configure the necessary software

tpaexec deploy clustername

# 4. Testing: make sure everything is working as expected

tpaexec test clustername
```

You can run TPA from your laptop, an EC2 instance, or any machine that can reach the cluster’s servers over the network.

Here’s a *TPA capabilities and supported software* .

### Configuration

The *Cluster configuration*

command generates a simple YAML configuration file to describe a cluster, based on the options you select. The configuration is ready for immediate use, but you can modify it to better suit your needs. Editing the configuration file is the usual way to *Cluster configuration* , both before and after it’s created.

At this stage, you must select an architecture and a platform for the cluster. An **architecture** is a recommended layout of servers and software to set up Postgres for a specific purpose. Examples include “M1” (Postgres with a primary and streaming replicas) and “PGD-Always-ON” (EDB Postgres Distributed 5 in an Always On configuration). A **platform** is a means to host the servers to deploy any architecture, e.g., AWS, Docker, or bare-metal servers.

### Provisioning

The *tpaexec provision*

command creates instances and other resources required by the cluster. The details of the process depend on the architecture (e.g., M1) and platform (e.g., AWS) that you selected while configuring the cluster.

For example, given AWS access with the necessary privileges, TPA will provision EC2 instances, VPCs, subnets, routing tables, internet gateways, security groups, EBS volumes, elastic IPs, etc.

You can also “provision” existing servers by selecting the “bare” platform and providing connection details. Whether these are bare metal servers or those provisioned separately on a cloud platform, they can be used just as if they had been created by TPA.

You are not restricted to a single platform—you can spread your cluster out across some AWS instances (in multiple regions) and some on-premise servers, or servers in other data centres, as needed.

At the end of the provisioning stage, you will have the required number of instances with the basic operating system installed, which TPA can access via SSH (with `sudo` to root).

## Deployment

The *tpaexec deploy*

command installs and configures Postgres and other software on the provisioned servers (which may or may not have been created by TPA; but it doesn’t matter who created them so long as SSH and `sudo` access is available). This includes setting up replication, backups, and so on.

At the end of the deployment stage, Postgres will be up and running.

## Testing

The *tpaexec test* command executes various architecture and platform-specific tests against the deployed cluster to ensure that it is working as expected.

At the end of the testing stage, you will have a fully-functioning cluster.

## Incremental changes

TPA is carefully designed so that provisioning, deployment, and testing are idempotent. You can run through them, make a change to `config.yml`, and run through the process again to deploy the change. If nothing has changed in the configuration or on the instances, then rerunning the entire process will not change anything either.

## Cluster management

Once your cluster is up and running, TPA provides convenient cluster management functions, including configuration changes, switchover, and zero-downtime minor-version upgrades. These features make it easier and safer to manage your cluster than making the changes by hand.

## Extensible through Ansible

TPA supports a *Instance configuration* , so you can do a lot just by editing `config.yml` and re-running `provision/deploy/test`. If you do need to go beyond what TPA already supports, you can write

- *TPA custom commands* , which make it simple to write playbooks to run on the cluster. Just create *commands/xyz.yml* in your cluster directory, and invoke it using *tpaexec xyz /path/to/cluster* . Ideal for any management tasks or processes that you need to automate.
- *TPA custom tests* , which augment the builtin tests with in-depth verifications specific to your environment and application. Using *tpaexec test* to run all tests in a uniform, repeatable way ensures that you will not miss out on anything important, either when dealing with a crisis, or just during routine cluster management.
- *TPA hooks* , which are invoked during various stages of the deployment. For example, tasks in *hooks/pre-deploy.yml* will be run before the main deployment; there are many other hooks, including *post-deploy* . This places the full range of Ansible functionality at your disposal.

### 1.1.3 It's just Postgres

TPA can create complex clusters with many features configured, but the result is just Postgres. The installation follows some conventions designed to make life simpler, but there is no hidden magic or anything standing in the way between you and the database. You can do everything on a TPA cluster that you could do on any other Postgres installation.

### 1.1.4 Versioning in TPA

TPA previously used a date-based versioning scheme whereby the major version was derived from the year. From version 23 TPA transitioned to semantic versioning, initially using a two-part `major-minor` pattern, then adopting full three-part semantic versioning in version 23.34.1. Under this scheme, the major version is only incremented where required to comply with the backward compatibility principle below.

#### Backwards compatibility

A key development principle of TPA is to maintain backwards compatibility so there is no reason for users to need anything other than the latest version of TPA. We define backwards compatibility as follows:

- A `config.yml` created with TPA X.a will be valid with TPA X.b where  $b \geq a$
- The cluster created from that `config.yml` will be maintainable and re-deployable with TPA X.b

Therefore, a new major version implies a break in backward compatibility. As such, we aim to avoid releasing major versions and will only do so in exceptional circumstances.

### 1.1.5 Getting started

Follow the *TPA installation* for your system, then *Cluster configuration* .

## 1.2 Release notes

The Trusted Postgres Architect documentation describes the latest version of Trusted Postgres Architect 23.

| Trusted Postgres Architect version                      | Release Date |
|---------------------------------------------------------|--------------|
| <i>Trusted Postgres Architect 23.41.0 release notes</i> | 26 Nov 2025  |
| <i>Trusted Postgres Architect 23.40.1 release notes</i> | 28 Oct 2025  |
| <i>Configuring lvm-direct for production</i>            | 02 Oct 2025  |
| <i>Trusted Postgres Architect 23.39.0 release notes</i> | 22 Aug 2025  |
| <i>Trusted Postgres Architect 23.38.1 release notes</i> | 25 Jun 2025  |
| <i>Trusted Postgres Architect 23.38.0 release notes</i> | 09 Jun 2025  |
| <i>Trusted Postgres Architect 23.37.0 release notes</i> | 24 Mar 2025  |
| <i>Trusted Postgres Architect 23.36.0 release notes</i> | 19 Feb 2025  |
| <i>Trusted Postgres Architect 23.35.0 release notes</i> | 25 Nov 2024  |
| <i>Trusted Postgres Architect 23.34.1 release notes</i> | 09 Sep 2024  |
| <i>Trusted Postgres Architect 23.34 release notes</i>   | 22 Aug 2024  |
| <i>Trusted Postgres Architect 23.33 release notes</i>   | 24 Jun 2024  |
| <i>Trusted Postgres Architect 23.32 release notes</i>   | 15 May 2024  |
| <i>Trusted Postgres Architect 23.31 release notes</i>   | 19 Mar 2024  |
| <i>Trusted Postgres Architect 23.30 release notes</i>   | 19 Mar 2024  |
| <i>Trusted Postgres Architect 23.29 release notes</i>   | 15 Feb 2024  |
| <i>Trusted Postgres Architect 23.28 release notes</i>   | 23 Jan 2024  |
| <i>Trusted Postgres Architect 23.27 release notes</i>   | 19 Dec 2023  |
| <i>Trusted Postgres Architect 23.26 release notes</i>   | 30 Nov 2023  |

continues on next page

Table 1 – continued from previous page

| Trusted Postgres Architect version                            | Release Date               |
|---------------------------------------------------------------|----------------------------|
| <i>Trusted Postgres Architect 23.25 release notes</i>         | 14 Nov 2023                |
| <i>Trusted Postgres Architect 23.24 release notes</i>         | 17 Oct 2023                |
| <i>Trusted Postgres Architect 23.23 release notes</i>         | 21 Sep 2023                |
| <i>Trusted Postgres Architect 23.22 release notes</i>         | 06 Sep 2023                |
| <i>Trusted Postgres Architect 23.21 release notes</i>         | 05 Sep 2023                |
| <i>Trusted Postgres Architect 23.20 release notes</i>         | 01 Aug 2023                |
| <i>Trusted Postgres Architect 23.19 release notes</i>         | 12 Jul 2023                |
| <i>Trusted Postgres Architect 23.18 release notes</i>         | 23 May 2023                |
| <i>Trusted Postgres Architect 23.17 release notes</i>         | 10 May 2023                |
| <i>Trusted Postgres Architect 23.16 release notes</i>         | 21 Mar 2023                |
| <i>Trusted Postgres Architect 23.15 release notes</i>         | 15 Mar 2023                |
| <i>Trusted Postgres Architect 23.14 release notes</i>         | 23 Feb 2023                |
| <i>Trusted Postgres Architect 23.13 release notes</i>         | 22 Feb 2023                |
| <i>Trusted Postgres Architect 23.12 release notes</i>         | 21 Feb 2023                |
| <i>Trusted Postgres Architect 23.1 to 23.11 release notes</i> | 21 Jun 2023 to 31 Jan 2023 |

## 1.2.1 Trusted Postgres Architect 23.41.0 release notes

Released: 26 November 2025

!!!Warning Known Issue When using TPA 23.35 or later to `deploy` to a PGD/BDR cluster created with an earlier version of TPA the `barman` user will have its superuser permissions revoked. This means that attempts by this user to connect to the `postgres` database will fail. This will cause the `Barman check` command to report a failure. The workaround for this issue is to restore these permissions with a post-deploy hook.

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.41.0 include the following:

### Highlights

- Perform updates of all cluster components - not just Postgres and PGD
- Hardening of the default Patroni configuration, with many new options to fine-tune security measures
- Support for minor version upgrades of PGD 6 (PGD-X and PGD-S)
- Support for Postgres 18

### Enhancements

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Addresses |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <b>Secure etcd and Patroni defaults for new <code>tpaexec configure</code> clusters.</b> When creating a new cluster configuration using <code>tpaexec configure</code> , the generated <code>config.yml</code> file will now automatically enable SSL/TLS encryption and mutual TLS (mTLS) authentication for both <code>etcd</code> and <code>Patroni</code> by default. This ensures that new deployments start with a secure communication baseline out-of-the-box. The following variables are now set by default for new configurations: • <code>etcd_ssl_enabled: true</code> • <code>etcd_authentication_mode: mtls</code> • <code>patroni_ssl_enabled: true</code> • <code>patroni_authentication_mode: mtls</code> Existing cluster configurations are unaffected by this change. |           |

continues on next page

Table 2 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Addresses |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <p><b>Support upgrade of components M1 and PGD/BDR clusters.</b> All cluster software components (failover managers, PEM, Barman, etc.) can now be updated, either to a specific version by specifying the component package version in <i>config.yml</i> or to the latest available when it is not provided. A list of components for upgrade can be passed to the <i>--components</i> flag as a comma-separated list. If this flag is not passed, <i>postgres</i> will be updated. Specifying <i>--components=all</i> will update all applicable software components.</p>                                                                                                                                                                                                                                                                                                                                                                                                                               |           |
| <p><b>Improved automation of PGD5.9+ to PGD6 upgrade.</b> For a PGD v5.9 to v6 migration, Connection Manager is a requirement instead of PGD proxy. A new command <i>switch2cm</i> has been introduced to facilitate this transition. The new command requires that the cluster's <i>config.yml</i> must be prepared by running <i>tpaexec reconfigure &lt;cluster&gt;; --enable-connection-manager</i>. <i>Upgrading from PGD-Always-ON to PGD-X</i></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |           |
| <p><b>Communication with the <i>etcd</i> cluster can now be secured using SSL/TLS.</b> A new configuration parameter, <i>etcd_ssl_enabled</i>, has been introduced to enable TLS encryption for all <i>etcd</i> communication. This significantly hardens the security of the cluster's distributed control plane. When <i>etcd_ssl_enabled</i> is set to <i>true</i>: • The deployment process automatically generates the required TLS certificates for each <i>etcd</i> node. These certificates now include both DNS and IP Subject Alternative Names (SANs) for robust validation. • Both peer-to-peer and client-server <i>etcd</i> connections are configured to use HTTPS. • Downstream components, including <i>Patroni</i> and internal health checks using <i>etcdctl</i>, are automatically configured to connect to <i>etcd</i> securely over HTTPS. To maintain backward compatibility with existing deployments, this feature is disabled by default (<i>etcd_ssl_enabled: false</i>).</p> |           |
| <p><b>Mutual TLS (mTLS) Authentication for Patroni REST API.</b> Support for mutual TLS has been added to provide a more secure, certificate-based authentication method for the Patroni REST API. A new configuration variable, <i>patroni_authentication_mode</i>, is now available. You can set this to <i>mtls</i> to enable this feature. When enabled (which requires <i>patroni_ssl_enabled: true</i>), the deployment automation will: • Configure the Patroni REST API server to require and validate client certificates. • Automatically configure clients, including <i>patronictl</i> and HAProxy health checks, with the necessary client certificates and keys to connect securely. The default authentication mode remains <i>basic</i> (username and password) to ensure backward compatibility with existing deployments.</p>                                                                                                                                                           |           |
| <p><b>Improve security of Patroni REST API with access <i>allowlist</i>.</b> The security of the Patroni REST API has been enhanced by configuring an <i>allowlist</i> for unsafe API endpoints (which accept <i>POST</i>, <i>PUT</i>, <i>PATCH</i>, <i>DELETE</i>). This measure prevents unauthorized nodes from performing administrative actions, such as restarting a database node or changing the cluster configuration. The <i>allowlist</i> is dynamically populated with the hostnames of the Patroni cluster members based on the Ansible inventory, ensuring that only authorized nodes can perform requests to these sensitive API endpoints.</p>                                                                                                                                                                                                                                                                                                                                            |           |
| <p><b>TPA now supports minor version upgrades for PGD-X.</b> The <i>tpaexec upgrade</i> command can now perform minor version upgrades of postgres and PGD on a cluster running the PGD-X architecture. The upgrade process checks cluster health and upgrades the nodes one at a time. This scenario supports the use of <i>update_hosts</i> variable to upgrade a subset of the cluster when used accordingly to the best practice recommendations found in the documentation.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |           |
| <p><b>TPA now supports minor version upgrades for PGD-S.</b> The <i>tpaexec upgrade</i> command can now perform minor version upgrades of postgres and PGD on a cluster running the PGD-S architecture. The upgrade process checks cluster health and upgrades the nodes one at a time. This scenario supports the use of <i>update_hosts</i> variable to upgrade a subset of the cluster when used accordingly to the best practice recommendations found in the documentation.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |           |

continues on next page

Table 2 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Addresses |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <p><b>Enable etcd Basic Authentication and Role-Based Access Control.</b> This release introduces a new capability to secure the etcd distributed key-value store with client authentication. You can now enable basic authentication and configure a dedicated user for Patroni with granular, least-privilege permissions. A new configuration variable, <i>etcd_authentication_mode</i>, has been added to control the feature. When set to <i>basic</i>, the deployment automation will:</p> <ul style="list-style-type: none"> <li>• Set up an administrative <i>root</i> user.</li> <li>• Create a dedicated etcd user for Patroni, given by the new <i>patroni_etcd_user</i> configuration option, and grant it read-write permissions limited to its cluster's key prefix (e.g., <i>/tpa/cluster_name</i>).</li> <li>• Configure Patroni to automatically use these credentials to securely connect and authenticate with etcd.</li> </ul> <p>To maintain backward compatibility, this feature is disabled by default (<i>etcd_authentication_mode: none</i>).</p> <p><b>Mutual TLS (mTLS) Authentication for `etcd`.</b> Support for mutual TLS has been added to secure all <i>etcd</i> communication. A new <i>mtls</i> option is now available for the <i>etcd_authentication_mode</i> variable. When <i>etcd_authentication_mode</i> is set to <i>mtls</i> (which requires <i>etcd_ssl_enabled: true</i>), the deployment automation will enforce certificate-based authentication for:</p> <ul style="list-style-type: none"> <li>• Peer-to-peer communication between all <i>etcd</i> nodes.</li> <li>• Client-server communication from clients like Patroni.</li> </ul> <p>TPA automatically configures both the <i>etcd</i> servers and the Patroni clients with the necessary certificates and keys to ensure a fully secure and validated connection.</p> <p><b>Introduced new instance-level variables to control how servers and agents are displayed in PEM.</b> Introduced four new instance variables for use with the <i>pem-agent</i> role: <i>monitoring_group</i> - <i>monitoring_agent_group</i> - <i>monitoring_cluster</i> - <i>monitoring_agent_cluster</i>. Also introduce the instance variable <i>pem_agent_group</i> for use with the <i>pem-server</i> role alongside the existing variable <i>pem-server-group</i>. Collectively, these variables determine which group and cluster the servers and agents are assigned to in PEM. See the documentation for a full explanation.</p> <p><b>Added support for validity period in OpenSSL generated certificates.</b> Previously, some SSL certificates (OpenVPN, PEM and TPA itself) used in TPA didn't come with a way of specifying the number of days before expiration. This change adds a variable called <i>openssl_certificate_validity</i> (default 3650 days), that can be used to determine the number of days an SSL certificate will last before expiring. In order to force TPA to reload a new certificate with a new expiration date, you must pass the appropriate option to <i>tpaexec deploy -e pem_web_server_renew_tls_certificates=true</i> or <i>tpaexec deploy -e tpa_tls_renew_tls_certificates=true</i>).</p> |           |

## Changes

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Addresses |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <p><b>Changed the default EFM password encryption to scram-sha-256.</b> TPA now defaults to <i>efm_user_password_encryption: scram-sha-256</i> for all new clusters configured with EFM as the failover manager. This change applies when running <i>tpaexec configure</i> with <i>--enable-efm</i> or <i>--failover-manager efm</i>. This addresses the deprecation of MD5 password encryption in PostgreSQL 18 and provides improved security for all PostgreSQL versions. Existing clusters are unaffected. Users can override this default by explicitly setting <i>efm_user_password_encryption: md5</i> in <i>config.yml</i> if needed, though MD5 support will be removed from TPA in a future release.</p> |           |

continues on next page

Table 3 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Addresses    |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| <b>The <code>tpaexec configure</code> command now adds <code>efm_version</code> to <code>config.yml</code>.</b> Support for explicitly adding the <code>efm_version</code> if introduced when executing <code>tpaexec configure</code> and choosing EFM as the failover manager. This can be seen when invoking <code>tpaexec</code> as follows: <code>tpaexec configure my-cluster-dir -a M1 --enable-efm --efm-version 5.2 (u2026)</code> . Additionally, documentation covering the <code>efm_version</code> selection has been added. |              |
| <b>TPA now runs <code>pgd-proxy</code> instead of <code>journalctl</code> for PGD Proxy version detection.</b> There are some edge cases where <code>journalctl</code> does not return the expected version. Hence, TPA now runs <code>pgd-proxy</code> directly to get the version information instead.                                                                                                                                                                                                                                  | 46753, 46426 |
| <b>Moved the pre-deploy hook outside repositories role.</b> In previous versions, this hook was part of the repositories role. When TPA introduced the possibility to exclude this role from deployments, it created a situation whereby this hook could be unintentionally skipped. To avoid this situation, we moved the hook outside the role so it's always executed.                                                                                                                                                                 | 52412        |
| <b>Added <code>auto.basebackup</code> property for EFM 5.2 and above.</b> Starting with EFM 5.2, there will be a new property <code>auto.basebackup</code> that adds on to the new auto-rebuild capability. Users can use the existing <code>pg_rewind</code> feature, <code>pg_basebackup</code> , or tell EFM to use both ( <code>pg_basebackup</code> will be run if <code>pg_rewind</code> fails).                                                                                                                                    |              |
| <b>Updated the internal dictionary of the AWS regions.</b> The change updates the internal dictionary of the AWS regions to expand the coverage of the supported regions and corresponding AZs.                                                                                                                                                                                                                                                                                                                                           |              |
| <b>Documented <code>postgres_port</code> in TPA documentation.</b> Added detail of the parameter <code>postgres_port</code> in the TPA documentation. <code>postgres_port</code> is the variable allowing the customization of the listening port for Postgres. This variable ensures consistent use of the set value across all components of the cluster relying on the information.                                                                                                                                                    |              |
| <b>Disabled TRACK and TRACE methods on the PEM HTTPD server.</b> PEM 10.3 contains a fix for HTTPD configuration, moving the directive to rewrite TRACK/TRACE requests to the <code>&lt;VirtualHost&gt;</code> section for the default virtual host on port defined by the <code>pem_server_ssl_port</code> variable. This change applies the same fix to all new deployments of PEM, regardless of version.                                                                                                                              | 54529        |
| <b>Added <code>efm-post-config</code> hook.</b> A new <code>efm-post-config</code> hook has been introduced to be able to execute new tasks after the <code>efm</code> configuration has taken place. New documentation explaining how to use it has also been developed under the 'TPA hooks' section.                                                                                                                                                                                                                                   |              |
| <b>Added support for pgaudit and postgis plugins on PG18.</b> Adding PG18 support for those plugins. The version of those depends on the Postgres version being supported as well. We have removed versions 11 and 12 from the plugin lists as they are no longer widely supported.                                                                                                                                                                                                                                                       |              |
| <b>Start PGD proxy after the network is available.</b> This service needs the network service to be up and running before it can be started. <code>Wants</code> and <code>After</code> directives have been added to the PGD Proxy service unit to ensure that.                                                                                                                                                                                                                                                                           | 54251, 52731 |
| <b>Install <code>libpq5</code> on PEM server v10+ for RHEL.</b> PEM 10.1 added <code>libpq5</code> as a dependency to fix <code>psycopg</code> errors related to <code>kerberos</code> authentication on RHEL 9 servers. To ensure <code>libpq5</code> is installed, TPA now explicitly installs it on RHEL servers when the PEM server version is 10 or greater.                                                                                                                                                                         |              |

continues on next page

Table 3 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Addresses |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <p><b>Install <code>'libcurl-full'</code> when required.</b> <i>libcurl-minimal</i> is installed certain RHEL-esque images, and while it satisfies the package dependencies for <i>edb-pem-agent</i>, it does NOT provide the complete set of APIs required (namely for SMTP) If it is installed, we remove it and install the <i>libcurl-full</i> meta-package to ensure full <i>libcurl</i> functionality exists for the PEM agent. Because the <i>edb-pem</i> package depends on the <i>edb-pem-agent</i> package, the <i>libcurl-minimal</i> package must also be removed on the PEM server in order to avoid dependency resolution conflicts. This is done by installing the package with the <i>allowerase</i> parameter.</p> |           |

## Bug Fixes

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Addresses |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <p><b>Patroni REST API SSL/TLS connections are now correctly configured and secure.</b> Previously, enabling SSL for the Patroni REST API (<i>patroni_ssl_enabled: true</i>) used a self-signed certificate. This resulted in <i>CERTIFICATE_VERIFY_FAILED</i> errors for clients like <i>patronictl</i>, making the feature unusable in a secure configuration. This has been fixed by implementing a proper certificate generation process. The Patroni REST API server is now configured with a server certificate signed by a Certificate Authority (CA) created by TPA for the TPA cluster. Additionally, the <i>ctl</i> section of the <i>patroni.yml</i> configuration has been corrected to ensure <i>patronictl</i> is properly configured to use the CA for server validation and to present its own client certificate and key for authentication (mTLS). As a result, enabling SSL for the Patroni REST API now provides a fully secure and verifiable communication channel for all clients.</p> |           |
| <p><b>Fixed race conditions during concurrent updates to the Barman <code>`.pgpass`</code> file.</b> Previously, when multiple backed-up nodes attempted to update the Barman node's <i>.pgpass</i> file concurrently, a race condition could occur, resulting in partial or incomplete entries. TPA now serializes this process, ensuring that the Barman node's <i>.pgpass</i> file is updated sequentially to prevent loss.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 52460     |
| <p><b>Fix <code>'unwanted_packages'</code> to support <code>'common'</code> key that applies to all distribution.</b> Previously, when setting <i>unwanted_packages</i> using the <i>common</i> key was not supported, which differs from the behavior of other package dictionary that can be set in <i>config.yml</i>. This fix introduces the support for the <i>common</i> key, ensuring any package listed under this key will be removed if present on any distribution in use in the cluster.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |           |
| <p><b>Fixed task selection for <i>pgdcli</i>.</b> Fixed a bug whereby the <i>pgdcli</i> role would be run on <i>bdr</i> nodes even when task selectors were used which should implicitly exclude it.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |           |
| <p><b>Fixed a fatal error during Docker container provisioning.</b> When creating containers on modern versions of docker using user-defined networks, <i>provision</i> would sometime crash with a 'dict object' has no attribute 'IPAddress' error. This was caused by an outdated method of discovering the container's IP address. The logic has been updated to be compatible with modern Docker network structures, making container creation more robust.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |           |
| <p><b>Fixed <code>'etcd'</code> package installation on SLES.</b> Deployments on SLES-based distributions that require <i>etcd</i> (such as those using Patroni) would previously fail because the <i>etcd</i> package could not be found in the standard repositories. This has been resolved by enabling the PGDG <i>extras</i> repository, which provides the necessary package and allows the installation to succeed.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |           |

continues on next page

Table 4 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Addresses |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <b>Fixed an issue whereby `tpaexec test` did not remove data after testing in PGD clusters.</b> After deploying a cluster and running <i>tpaexec test</i> on that cluster, creates an additional entry in <i>bdr.replication_set_table</i> which is not later removed. The fix makes the test run on a temporary schema called <i>tpatest</i> , which is later removed after the test is finished.                                                                                                                                                  | 53224     |
| <b>Fixed issue with missing PGD CLI on proxy nodes when performing upgrade (3.7 to 5).</b> When performing a regular deploy, or an upgrade from a 'BDR-Always-ON' setup to 'PGD-Always-ON', the proxy nodes (previously deployed with the 'harp-proxy' role), will not come shipped with the PGD CLI when trying to install it with the 'pgd-proxy' role, necessary to perform the upgrade. This fix ensures that the PGD CLI will be installed on the proxy nodes and nodes with a 'bdr' role, also on the nodes featuring only the 'pgdeli' role. |           |

## 1.2.2 Trusted Postgres Architect 23.40.1 release notes

Released: 28 October 2025

This is a bug fixes release. The issues resolved in Trusted Postgres Architect 23.40.01 include the following:

### Bug Fixes

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Addresses |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <b>Fixed an issue whereby the PGD `alter_node_option` task would run for physical replicas.</b> In a <i>PGD-AlwaysON</i> cluster with a BDR node that is also a physical replica, the upgrade process would fail during the <i>alter_node_option</i> task because a BDR node is expected to have RAFT enabled, but the physical replica BDR node does not. TPA now skips this task on any node with the <i>replica</i> role, allowing upgrade to complete successfully.                                                                                                                                                                                                                                                                                                                                                      | 51997     |
| <b>Fixed unnecessary output after deploy by removing the post-deploy check for <i>pgbouncer_auth_function</i> in <i>pg_catalog</i>.</b> TPA recently relocated the <i>pgbouncer_auth_function</i> used to authenticate users when <i>pgbouncer</i> is in use. This was done in an effort to better follow the best practices and avoid issues during upgrade where <i>pg_catalog</i> function would not be part of the dump of the database. TPA introduced a post deploy check to help users still relying on the auth function defined in <i>pg_catalog</i> to identify the problem and make sure that they are able to fix it before removing the function from their systems. This check was used as a transition mechanism and is now being removed. This resolves issues with unwanted output generated by that check. | 100065    |
| <b>Fixed a misconfiguration of the <i>.pgpass</i> file when <i>backup_name</i> is used.</b> Previously, when using the <i>backup_name</i> variable, TPA would incorrectly use this name in <i>pgpass</i> for the Barman server. This has been fixed and the documentation has been updated.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |           |
| <b>TPA now ensures EDB repos are used when Beacon Agent is included in a cluster.</b> TPA can deploy the Beacon Agent to monitor the health of a system. However, it needs access to a valid EDB repository. This change ensures that the Standard repository is enabled when the Beacon Agent is to be deployed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |           |
| <b>Fixed issue whereby Barman could not configure the <i>log_file</i> parameter correctly.</b> When setting up a barman node, the variable ' <i>barman_log_file</i> ' was not setting up an alternative directory to place the log file. Instead, it used the default location ( <i>/var/log/barman/barman.log</i> ). Bug and fix reported by voxnyx ( <a href="https://github.com/voxnyx">https://github.com/voxnyx</a> ).                                                                                                                                                                                                                                                                                                                                                                                                  |           |

### 1.2.3 Trusted Postgres Architect 23.40.0 release notes

Released: 2 October 2025

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.40.0 include the following:

#### Enhancements

| Description                                                                                      | Addresses |
|--------------------------------------------------------------------------------------------------|-----------|
| Added support for upgrading EDB Postgres Distributed (PGD) v5 clusters to v6.                    |           |
| Extended support for configuring PGD-S clusters.                                                 |           |
| Introduced a <i>pgd-proxy-config</i> hook.                                                       | 49911     |
| Added three <i>release.vip.*</i> properties for EFM 5.1 and above.                               |           |
| TPA now allows the RSA key size to be set for certificates.                                      |           |
| Documented how TPA generates Postgres server TLS certificates and keys and how to override this. | 52049     |

#### Changes

| Description                                                     | Addresses |
|-----------------------------------------------------------------|-----------|
| TPA now supports EPAS clusters with repmgr as failover manager. |           |

#### Bug Fixes

| Description                                                                           | Addresses    |
|---------------------------------------------------------------------------------------|--------------|
| Fixed an issue whereby backup wasn't configured in PGD 6 clusters.                    |              |
| Fixed support for <i>update_host</i> variable on PGD-Always-ON.                       | 52558, 45181 |
| HAProxy health checks now correctly use HTTPS for SSL-enabled Patroni clusters.       |              |
| Fixed an issue with <i>pgaudit</i> extension when used with <i>edbpgge</i> Postgres.  |              |
| Fixed an issue where global proxy routing was not set up correctly in PGD 5.9.        |              |
| TPA will now correctly skip the <i>postgres/pgpass</i> task on PEM-agent Barman nodes |              |

### 1.2.4 Trusted Postgres Architect 23.39.0 release notes

Released: 22 August 2025

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.39.0 include the following:

#### Highlights

- Support for s390x architectures
- Support for *edb\_stat\_monitor* extension
- EFM property support for *auto.rewind*, *jdbc.properties*, *jdbc.loglevel*, and *check.vip.timeout*

#### Enhancements

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Addresses |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Added support for PEM on Ubuntu 24.04. There was no specific package defined for this distro which caused deployment errors. Additionally, this ensures that either <code>mod_wsgi</code> or <code>edb's mod_wsgi</code> module is enabled by default where it applies.                                                                                                                                                                                                                                                                                                                                                              |           |
| Added <code>edb_stat_monitor</code> to recognized extensions. When a user specifies <code>edb_stat_monitor</code> as an entry in either <code>postgres_extensions</code> or the list of extensions named under <code>postgres_databases</code> , TPA will handle installing the correct package, creating the extension and including it in the <code>shared_preload_libraries</code> .                                                                                                                                                                                                                                              |           |
| Add support for s390x (aka IBM Z and LinuxONE). TPA now supports the s390x CPU architecture. Supported s390x operating systems are RHEL 8 and 9 and SLES 15. TPA can run on these systems and use them as target hosts for cluster deployment. Because s390x binaries are not available for all PyPI packages, it is highly recommended that you install the <code>tpaexec-deps</code> package from EDB rather than rely on installation from PyPI during setup. Similarly, because there is not a PGDG RPM repo for s390x you must have access to EDB Repos for TPA to install Postgres and other cluster components from packages. |           |
| Added <code>auto.rewind</code> property for EFM 5.1 and above. Starting with EFM 5.1, there will be a new property <code>auto.rewind</code> that is part of a new feature to attempt to rebuild failed primary db servers.                                                                                                                                                                                                                                                                                                                                                                                                           |           |
| Add EFM <code>jdbc.properties</code> property support. This adds support for the <code>jdbc.properties</code> property in EFM 5.0 and later.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |           |
| Add <code>jdbc.loglevel</code> property for EFM 5.1 and above. Starting with EFM 5.1, there will be a new property <code>jdbc.loglevel</code> that is used to increase information logged from the JDBC driver. This property can be used to get more information when there are connection problems, e.g. when using <code>ssl</code> for database connections.                                                                                                                                                                                                                                                                     |           |
| Add <code>check.vip.timeout</code> property for EFM 5.1 and above. Starting with EFM 5.1, there will be a new property <code>check.vip.timeout</code> that is used to control how long EFM will keep checking is the VIP (if used) is reachable before promoting a standby.                                                                                                                                                                                                                                                                                                                                                          |           |

## Changes

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Addresses |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Ubuntu 20.04 (Focal) is now a legacy distribution. Ubuntu 20.04 is now out of support upstream and is no longer fully supported by TPA.                                                                                                                                                                                                                                                                                                                                                                                          |           |
| Removed support for obsolete EDB/2q repositories. Now that the old 2q repositories and EDB Repos 1.0 are no longer available, TPA will not try to use them or check for configuration related to them. All packages that were formerly available in those repositories are now available in the EDB Repos 2.0 repositories.                                                                                                                                                                                                      |           |
| TPA will now suppress PGDG repos when using EDB repos on SLES. When running <code>tpaexec configure</code> for a SLES system, use either EDB repositories or PGDG repositories, not both, matching the behavior of other operating systems. This ensures that packages such as <code>barman</code> , which are available in both places, are consistently sourced from the same repository as other packages.                                                                                                                    |           |
| TPA will now use <code>bdr_package_version</code> as default PGD Proxy version. <code>pgd_proxy_package_version</code> and <code>pgdcli_package_version</code> may be explicitly defined in <code>config.yml</code> . If they are not explicitly defined, they now default to the value of <code>bdr_package_version</code> , if it is set. If none of them is set, the latest available versions of packages are installed. Setting these package versions to different values is not supported and is only useful for testing. | 49150     |
| Reduced default for <code>vacuum_cost_delay</code> to 2ms. In recent Postgres versions, the suggested default for <code>vacuum_cost_delay</code> has been reduced to 2ms. TPA now matches this.                                                                                                                                                                                                                                                                                                                                  |           |

continues on next page

Table 10 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Addresses |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Documented node promotability logic and the <i>efm-not-promotable</i> role for EFM. Added a new section to the <i>efm.md</i> documentation that explains how TPA determines whether a node is eligible for promotion during failover. The update clarifies the rules for promotability, including the roles of witness nodes, cascading standbys, and nodes explicitly marked with the <i>efm-not-promotable</i> role. This enhancement should help users understand and control failover behavior, reducing the risk of unintended promotions in EFM-managed clusters. |           |
| Document <i>include_vars</i> behavior for templated variables. Since the <i>include_vars</i> module immediately parses and evaluates expressions, nested variables do not exist at the point they are loaded from <i>config.yml</i> and thus are undefined when evaluated in templated expressions. This is now documented with an example so it is clearer for users.                                                                                                                                                                                                  | 45722     |
| Use Docker images from Rocky Linux organization. TPA now uses the Rocky Linux organization's docker images rather than docker hub's "official" ones. These are more frequently updated and hence less likely to cause dependency problems with newer packages.                                                                                                                                                                                                                                                                                                          |           |
| Updated architecture info for PGD 6 architectures. The PGD6 architectures now have correct metadata and therefore appear as expected in the output of <i>tpaexec info architectures</i> .                                                                                                                                                                                                                                                                                                                                                                               |           |

## Bug Fixes

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Addresses |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Fixed an issue whereby <i>replication</i> user was missing from <i>.pgpass</i> on the primary node. The <i>replication_user</i> entry is required the <i>.pgpass</i> file for both <i>replica</i> and <i>primary</i> nodes managed by either <i>efm</i> or <i>patroni</i> as the <i>failover_manager</i> . Previously, it was only added to the <i>.pgpass</i> file for <i>replica</i> nodes, resulting in connection issues to the primary node after a switchover.                                                                                                                                                                                                                                                                                                                        | 48224     |
| Fixed an issue whereby TPA would attempt to enable routing for subscriber-only groups in PGD6 clusters. Previously, TPA would mistakenly try to enable routing for subscriber-only groups during PGD6 deployments. This property is not editable anymore as it forms core part of the definition of subscriber-only groups in PGD6.                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 49673     |
| Fixed an issue whereby PEM deployment would fail with PEM 10.1.1. The TPA task <i>Register PEM backend database server for monitoring and configuration</i> explicitly calls the <i>pem.setup</i> SQL function. This is not considered part of the public API of PEM and the signature changed between 10.1.0 and 10.1.1 causing TPA's call to this function to fail. This fix addresses the issue by adjusting the function call according to the installed PEM server version. To accommodate this change we have introduced a new <i>pem/server/facts</i> Ansible role that is responsible for collecting facts about the installed PEM version. This also means that TPA will no longer attempt to run <i>postgresexpert.sql</i> on PEM 10, where Postgres Expert is no longer present. |           |
| Fixed an issue whereby <i>harp_dcs_user</i> was not created when using multiple failover managers                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 49150     |
| Fixed an issue whereby <i>bdr_node_groups</i> could not be overridden. If a file supplied in the <i>--overrides-from</i> argument to <i>tpaexec configure</i> sets <i>bdr_node_groups</i> in <i>cluster_vars</i> , the contents of this will now be added to the node groups automatically created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 50550     |

continues on next page

Table 11 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Addresses |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Fixed issue whereby TPA would use <i>pgdproxy</i> user in DSN on PGD6 nodes. Since PGD version 6 has a built-in Connection Manager which replaces PGD Proxy, the <i>pgdproxy</i> Postgres user should not be used in any DSNs. A new variable <i>bdr_connection_manager_route_dsn</i> is introduced for defining connection strings to Connection Manager. Users can additionally define <i>bdr_connection_manager_dsn_attributes</i> in their <i>config.yml</i> to specify additional connection parameters in the DSN.                                                                                                                                                                                                                                                                          |           |
| Fixed an issue with <i>ignore_slots</i> for Patroni clusters with multiple Barman servers. For Patroni clusters TPA will now dynamically generate the <i>ignore_slots</i> setting corresponding to Barman nodes in the cluster. This change ensures that Patroni will not remove the physical slots created for each of the Barman nodes. This is especially important for the Patroni clusters with more than one backup servers; including shared Barman nodes for example.                                                                                                                                                                                                                                                                                                                     | 50748     |
| Fixed usability and improve reliability of the TPA Docker image. Upgraded Base Image to Debian Trixie to meet TPA's Python 3.12 minimum requirement, ensuring compatibility with modern dependencies. - Removed <i>--use-community-ansible</i> flag as it is no longer supported. - Included the <i>openssh-client</i> package to prevent <i>tpaexec</i> failures due to missing <i>ssh-keygen</i> , enhancing out-of-the-box functionality. - Introduced Build-Time Versioning: Added an <i>ARG</i> variable to pass the output of <i>git describe</i> from the host during build, storing it as <i>/opt/EDB/TPA/VERSION</i> for improved version tracking. - Revised <i>docker/README.md</i> and <i>docs/src/INSTALL-docker.md</i> to reflect the latest setup instructions and best practices. |           |
| Fixed an issue whereby Barman could not run <i>switch-wal</i> . For Postgres 15+, the Barman user is now created with the <i>pg_checkpoint</i> role. This allows Barman to run CHECKPOINT without the need of being superuser, <i>barman switch-wal --force</i> will not fail any more.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |           |

## 1.2.5 Trusted Postgres Architect 23.38.1 release notes

Released: 25 June 2025

Trusted Postgres Architect 23.38.1 is a bug-fix release that resolves the following issues:

### Bug Fixes

| Description                                                                                                                                                                                                                                                                                                                                                                                                              | Addresses |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Fixed an issue with the configure command for PGD6 using bare platform. Fixed a bug with <i>tpaexec</i> configure command on the newly released architecture for PGD6 (PGD-S and PGD-X) whereby trying to generate a cluster using <i>--platform bare</i> would result in an Unknown Platform error. This fix ensures that the configure command successfully generates a PGD6 configuration file for the bare platform. | 49673     |

continues on next page

Table 12 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Addresses |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Fixed an issue whereby <code>pgpass_users</code> were not correctly added to the <code>.pgpass</code> file. Previously, any user in the <code>postgres_users</code> list specified with <code>generate_password: true</code> AND included in the <code>pgpass_users</code> list would NOT be added to the <code>~postgres/.pgpass</code> file on the initial deploy because the user's password did not yet exist when the <code>pgpass</code> task was executed, thus the user was skipped. This is fixed by invoking the <code>pgpass</code> task once more after all the Postgres users have been created. The <code>repmgr</code> and <code>replication</code> users were previously included in the <code>default_pgpass_users</code> list. They are now added to <code>pgpass_users</code> and hence to the <code>.pgpass</code> file as required by the replication manager in use. The <code>postgres_user</code> ( <code>postgres</code> or <code>enterprisedb</code> by default) is still part of the <code>default_pgpass_users</code> list. If this is overridden by a <code>pgpass_users</code> list in <code>config.yml</code> that does NOT include <code>postgres_user</code> , a PEM-enabled cluster will fail to register agents as it needs the encrypted <code>postgres_user</code> password. This is fixed by adding the <code>postgres_user</code> to the <code>~postgres/.pgpass</code> file as part of the PEM agent tasks. |           |
| Fixed an issue whereby an invalid <code>cluster_vars</code> dictionary would lead to a failed deploy or an unexpected configuration. Previously, if an invalid dictionary is set as the <code>cluster_vars</code> dictionary in <code>config.yml</code> (such as <code>cluster_vars</code> variables referencing other <code>cluster_vars</code> variables), TPA would swallow any Ansible errors by falling back to the default value of an empty dictionary. This resulted in every <code>cluster_vars</code> variable being undefined, so each was set to its TPA-default value. The resulting cluster would be entirely different than what the user specified in their <code>config.yml</code> file. This is fixed by asserting that the <code>cluster_vars</code> dictionary is defined and non-empty when the configuration file is loaded. Also as a final bailout, the <code>cluster_vars</code> variable now no longer defaults to an empty dictionary. This allows Ansible to throw an error when creating <code>group_vars</code> and terminate.                                                                                                                                                                                                                                                                                                                                                                                        |           |

## 1.2.6 Trusted Postgres Architect 23.38.0 release notes

Released: 9 June 2025

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.38.0 include the following:

### Highlights

- Support for deploying EDB Postgres Distributed version 6 in both Expanded and Essential architectures
- Improvements to PgBouncer implementation including a mitigation for CVE-2025-2291
- Support for PEM 10.1

### Enhancements

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Addresses |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Support for PGD6 architectures. TPA can now configure and deploy clusters using the PGD-X and PGD-S architectures based on PGD6. The PGD-S architecture implements PGD Essential and the PGD-X architecture implements PGD Expanded. These architectures have sensible default configurations and also accept various configure options to customize their behavior. PGD 6 deployments no longer include <code>pgd-proxy</code> ; instead, PGD's built-in Connection Manager is configured. Testing support for the new architectures is added. |           |

continues on next page

Table 13 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Addresses |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Configure PEM to monitor Barman when both are present in a cluster. When a cluster is configured with PEM enabled (using the <code>--enable-pem</code> option) and includes a Barman node, the following actions are now performed automatically: - <code>enable_pg_backup_api</code> is set to <code>true</code> in <code>config.yml</code> - The <code>pem-agent</code> role is assigned to the Barman node - The Barman endpoint is registered with the local PEM agent. These changes simplify setup and ensure seamless integration between PEM and Barman. |           |
| Added support for Rocky Linux 9 on AWS. TPA now supports configuring a cluster using Rocky Linux 9.5 on the <code>aws</code> platform. This is now the default version for Rocky Linux on AWS if a version is not specified.                                                                                                                                                                                                                                                                                                                                     |           |
| Added support for <code>pg_backup_api</code> on SLES 15. TPA will now configure <code>pg_backup_api</code> is on SUSE Linux Enterprise Server 15 (SLES 15) when PEM monitoring is enabled and a Barman node is present in the cluster.                                                                                                                                                                                                                                                                                                                           |           |

## Changes

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Addresses   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| Treat <code>PEM_DB_PORT</code> as a string for PEM 10.1 and above. PEM 10.1 adds support for multi-host connection strings from the web application to the backend servers. To support this change, the <code>PEM_DB_PORT</code> parameter in PEM's <code>config_setup.py</code> file is now a string rather than an integer. While TPA does not yet support deploying HA PEM configurations, TPA will now correctly set this parameter as a string when the PEM version is 10.1 or greater.                                                                                                                                                          |             |
| TPA will now skip repository checks when <code>repo</code> is excluded from tasks. The <code>repo</code> tag is available for exclusion, but previously would only skip tasks under the <code>sys/repositories</code> role. Now it also skips over the initialization tasks which check which repositories to use and the verifies the credentials to access them are provided.                                                                                                                                                                                                                                                                       |             |
| Improved the behavior of <code>postgres_package_version</code> . Setting <code>postgres_package_version</code> will now cause TPA to install the selected version of various postgres-related components on Debian or Ubuntu systems installing EDB Postgres Advanced Server or EDB Postgres Extended Server. This avoids dependency resolution problems when newer package versions are visible in repositories.                                                                                                                                                                                                                                     |             |
| Added a new task selectors <code>create_postgres_system_user</code> and <code>create_pgproxy_system_user</code> . Added new task selectors that allow to skip the <code>postgres_user</code> and <code>pgd_proxy_user</code> operating system user. This allows clusters to use remote users created by a centralized user management such as NIS. This can be set in <code>config.yml</code> : <code>cluster_vars:</code><br><code>excluded_tasks:</code> - <code>create_postgres_system_user</code> - <code>create_pgproxy_system_user</code>                                                                                                       | 48601,44388 |
| TPA will now redirect PgBouncer to the new primary in M1 <code>repmgr</code> clusters during switchover. TPA will now ensure that PgBouncer instances are redirected to the new primary node after using the <code>switchover</code> command in a <code>repmgr</code> + PgBouncer cluster that has <code>repmgr_redirect_pg_bouncer</code> set to <code>true</code> . The <code>tpaexec switchover</code> command will now ensure that PgBouncer instance connect to the new primary node. A new <code>revert_redirect</code> variable can also be set as extra-variable after a first switchover is done to revert back to the initial primary node. |             |
| Made <code>key_id/gpgkey</code> optional in custom repository definitions. The <code>key_id</code> and <code>gpgkey</code> parameters (for <code>apt</code> and <code>yum</code> custom repositories definition) are not required by the underlying modules, there are use cases where this is not easy to provide ahead of installation. With this change, TPA does not make it mandatory to provide those in custom repository definitions.                                                                                                                                                                                                         |             |

continues on next page

Table 14 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Addresses |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Fixed <i>verify-settings</i> check in <i>tpaexec test</i> for PGD CLI 5.7.0+. The output for the PGD CLI command <i>pgd verify-settings</i> changed in PGD 5.7.0. TPA now correctly parses the output when using version above 5.7.0 of PGD CLI. Note, <i>verify-settings</i> will be deprecated along with other commands in future PGD releases. those commands are now wrapper calling the new commands until the deprecation occurs.                                                                                                                                                                    |           |
| Improved logic for granting permissions to <i>barman_role</i> . TPA now uses the <i>postgresql_privs</i> module to apply the grant on <i>barman_role</i> , so that changes are only applied when needed. In addition, in PGD clusters, TPA will use the <i>bdr_database</i> on second deploys so the DDL is replicated across the cluster by PGD.                                                                                                                                                                                                                                                           |           |
| Documented <i>cluster_vars</i> variable templating in <i>config.yml</i> . Added documentation to explain correct templating procedure for variables defined under <i>cluster_vars</i> with a worked example in order to avoid confusion from unexpected behavior associated with inventory variables not being defined when improperly templated in <i>config.yml</i> .                                                                                                                                                                                                                                     | 48797     |
| TPA will now raise an <i>ArchitectureError</i> for if an invalid CIDR is passed to <i>--network</i> flag during configure. By default, the Python standard library <i>ipaddress</i> package enforces 'strict' interpretation of the CIDR, whereby the IP used should be the network address of the range. Previously, any IP passed to the <i>--network</i> flag that contained host bits would dump a stacktrace due to the raised <i>ValueError</i> . That exception is now caught and an <i>ArchitectureError</i> is raised to display a clear message to the user about the <i>--network</i> parameter. |           |
| Added a default value for EFM <i>application.name</i> property. If the EFM <i>application.name</i> property is not set for a node, TPA will use the Postgres <i>cluster_name</i> property as a default. EFM uses this value when performing a switchover or when building a new standby database.                                                                                                                                                                                                                                                                                                           |           |
| TPA now uses the EDB repository setup script on SUSE. Previously, TPA did not use the EDB repository setup script on SUSE because it did not work on repeat deploys. Zypper would raise because the repositories that the script attempts to install already exist, and require unique names. Now that the repository setup script task is skipped if the repositories are already installed, this issue is not encountered.                                                                                                                                                                                |           |

## Bug Fixes

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Addresses    |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| <p>TPA will now create the <code>pgbouncer_get_auth()</code> function in dedicated database. The <code>pgbouncer_get_auth()</code> function was created in the <code>pg_catalog</code> schema and execute granted to the <code>pgbouncer_auth_user</code>. This function was created in every database, but this was not necessary for PgBouncer. A failure may be encountered during the <i>pgd node upgrade</i> process when this function was created in the <code>pg_catalog</code> schema as it is not included in the dump created by <code>pg_dump</code>. A later task attempts to run a <code>GRANT</code> on this function and fails, as the function is not restored since it was not originally dumped. Now this function is only created in a single database, named under the <code>pgbouncer_auth_database</code> variable in <code>config.yml</code>, which defaults to <code>pgbouncer_auth_database</code> if not included. It is only created if at least one instance with <code>pgbouncer</code> role is included in the cluster. A warning is also issued during deploy and upgrade if any databases define this function under the <code>pg_catalog</code> schema, as a future TPA release may remove the function from that schema. The <code>pgbouncer_get_auth()</code> function itself used by PgBouncer <code>auth_query</code> has been updated to address CVE-2025-2291. This vulnerability allowed for authentication using expired passwords, a <code>_query</code> mechanism did not consider the <code>VALID UNTIL</code> attribute set in PostgreSQL for user passwords.</p> <p>Fixed an issue whereby some tasks were incorrectly skipped when the <code>--check</code> option was used. In PGD clusters without HARP, the <i>Read current configuration file if exists</i> task needs to run in check mode to ensure we have the information available to correctly skip the following HARP check task. However, by default Ansible skips tasks using the <code>shell</code> module during check mode, meaning this task did not run, resulting in a spurious failure on subsequent tasks. We now let Ansible know that this task has to be run.</p> <p>Fixed a bug whereby settings added to <code>ignore_slots</code> via <code>cluster_vars['patroni_conf_settings']['bootstrap']['dcs']</code> were not merged into the eventual config.</p> | 42911, 45068 |
| <p>Fixed an issue where RAFT checks for BDR nodes with replica role were not skipped during upgrade. Physical replication of a <i>subscriber-only</i> node can be achieved in a <i>PGD</i> cluster by installing <code>repmgr</code> as a failover-manager and designating the <i>subscriber-only</i> node as the <i>primary</i> and listing another BDR data node as the <i>backup</i>; this backup node is given the <i>replica</i> role. This configuration would result in the PGD upgrade process failing, since TPA expects BDR data nodes to have RAFT enabled, but the physical replica BDR data node (with both <i>replica</i> and <i>bdr</i> roles) by design does not. As a fix, certain BDR-specific tasks in the upgrade process now skip any node that has a <i>replica</i> role, allowing for a successful upgrade.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 46186        |

## 1.2.7 Trusted Postgres Architect 23.37.0 release notes

Released: 24 March 2025

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.37.0 include the following:

### Highlights

- The `upgrade` command now supports minor-version upgrades of Postgres in Patroni and EFM clusters
- Upgrades in `repmgr` clusters have also been improved

### Enhancements

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Addresses |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Minor-version Postgres upgrades for M1 + Patroni clusters. TPA can now upgrade Postgres to the latest minor version on an M1 cluster which uses Patroni as the failover manager. The upgrade process stops Barman on any Barman server in the cluster, then upgrades the replicas in the cluster. Then it switches to one replica as a temporary primary, upgrades Postgres on the original primary, and switches back to the original primary. Patroni's handling of the cluster is paused during the process and resumed afterwards. Then Barman is restarted and cluster health checks are run. | 102662    |
| Improved minor-version Postgres upgrade for M1 + repmgr clusters. Witness nodes were previously omitted from upgrade, they are now upgraded along with replicas. Postgres service restart is now more reliable and always run right after the package upgrade on the node is finished.                                                                                                                                                                                                                                                                                                             |           |
| Minor-version Postgres upgrades for M1 + EFM clusters. TPA can now upgrade Postgres to the latest minor version on an M1 cluster which uses EFM as the failover manager. The upgrade process stops Barman on any Barman server in the cluster, then upgrades the replicas in the clusters. Then it switches to one replica as a temporary primary, upgrades Postgres on the original primary, and switches back to the original primary. The EFM agent is started and stopped on the different servers at the correct times. Then Barman is restarted and cluster health checks are run.           | 103595    |
| Separated changed from unchanged tasks in output. In TPA's default output plugin, tasks which return "ok" but with no changes are now separated from ones that have reported changes, which are now highlighted in yellow.                                                                                                                                                                                                                                                                                                                                                                         |           |
| Improved the speed of Docker instance deprovisioning. When deprovisioning docker instances, TPA now kills the container instead of stopping it, and does so to all the instances in parallel.                                                                                                                                                                                                                                                                                                                                                                                                      |           |
| Added <i>check.num.sync.period</i> property for EFM 5.x. Starting with EFM 5.0, there is a new property <i>check.num.sync.period</i> that defines how often a primary agent will check to see if <i>num_sync</i> needs to change on the primary database. This can now be specified in <i>config.yml</i> .                                                                                                                                                                                                                                                                                         |           |

## Changes

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Addresses |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| TPA will now copy EFM config files if they are removed, even if no configuration changes. If either the <i>efm.nodes</i> or <i>efm.properties</i> configuration files do not exist in the top-level EFM directory, the <i>efm upgrade-conf</i> command copies them from the <i>/raw</i> directory, even if there have been no configuration changes. This amends previous behavior that required a configuration change before the <i>upgrade-conf</i> command would run and copy files.                                                                                                                                                                                                                             |           |
| TPA will now verify that the URI for EDB repository setup is accessible. The EDB repos are set up using the setup script following the EDB Repos 2.0 documentation: piping the cURL output to bash for execution. However, if a user passes a nonexistent <i>EDB_SUBSCRIPTION_TOKEN</i> or repository to cURL, the exit code gets silently swallowed and replaced with a 0 because bash executes an empty input. In this version, a request is dispatched to a repository's GPG key endpoint to ensure a <i>404</i> response is not returned <i>before</i> continuing to download the setup script. Additionally, tasks related to EDB repository set up are now skipped if the repository has already been set up.' |           |

continues on next page

Table 17 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                    | Addresses |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| TPA will now raise an <code>ArchitectureError</code> when BDR-Always-ON is configured with BDR version 5. An architecture error is now raised during <i>tpaexec configure</i> if <i>--bdr-version 5</i> is passed with <i>-a BDR-Always-ON</i> alerting the user that BDR version 5 should be used with <i>PGD-Always-ON</i> . |           |

## Bug Fixes

| Description                                                                                                                                                                                                                                                                                                                                             | Addresses |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Fixed <i>shared_preload_libraries</i> computation during deploy. Fixed a limitation of Ansible's handling of list ordering that would trigger unneeded and uncontrolled rewriting of the <i>shared_preload_libraries</i> and subsequently require a Postgres service restart, even on second deployment scenarios with no changes to the configuration. |           |
| Fixed an issue whereby unnecessary replication slots were created when configuring Patroni. When setting up a Patroni cluster, a replication slot was created for each etcd-only node. This caused problems because the unused slots cause the WAL to accumulate. Slots are now only created for the DB servers.                                        |           |
| Fixed <i>patronictl switchover</i> command usage. TPA will now correctly use <i>--leader</i> instead of the deprecated <i>--master</i> parameter when using <i>patronictl switchover</i> command.                                                                                                                                                       |           |
| Fixed an issue with counting instances in BDR-Always-ON clusters containing PEM. Fixed a bug whereby in certain circumstances, TPA would incorrectly calculate the number of instances in a BDR-Always-ON cluster with a PEM server, causing "tpaexec configure" to fail with "StopIteration".                                                          |           |
| Fixed an issue whereby a PEM server could be incorrectly picked as a cluster primary in EFM cluster. If a cluster is created with a pem-server, that backend is not monitored by EFM, hence, that node shouldn't be included when discovering a Postgres primary for the entire cluster. This behavior is now correctly implemented by TPA.             | 45279     |
| Fixed duplicated lines in .pgpass files. Fixed a bug whereby extra lines could be added to .pgpass for the same user when re-running 'tpaexec deploy'.                                                                                                                                                                                                  |           |

## 1.2.8 Trusted Postgres Architect 23.36.0 release notes

Released: 19 February 2025

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.36.0 include the following:

### Highlights

- Support for Ubuntu 24.04 and SLES 15 SP6
- Version specifiers for all cluster components
- Better support for EDB Failover Manager (EFM) 5

### Enhancements

| Description                                                                                                                                                                                                                                                                                                                                                                                                                            | Addresses |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Added support for package version specifiers for all cluster components. The following software packages now accept an <code>--xxx-package-version</code> option to the <code>tpaexec configure</code> command, which populates <code>xxx_package_version</code> in the generated <code>config.yml</code> - barman - pgbouncer - beacon-agent - etcd - patroni - pem-server - pem-agent - pg_backup_api - pgd_proxy - pgdcli - repmgr. |           |
| Added support for Ubuntu Noble 24.04. TPA runs on Ubuntu 24.04 and supports this OS as a target host. <b>EDB packages for Ubuntu 24.04 will be generally available at the end of March 2025. Before that, deployments needing EDB packages will likely fail.</b>                                                                                                                                                                       |           |
| Added support for EFM 5 "auto resume" properties. Starting with EFM 5.0, the <code>auto.resume.period</code> property has been broken into two properties, one for the startup case and one for the db failure case. This change adds the correct properties based on the <code>efm_version</code> being used.                                                                                                                         |           |
| Added support for EFM 5 'backup.wal' property. Starting with EFM 5.0, a new property 'backup.wal' has been added. This change adds the new property if the version of EFM is 5 or higher.                                                                                                                                                                                                                                              |           |
| Added support for SLES 15 SP6. When SLES is requested at configure-time, TPA will now install SLES 15 SP6. The Docker and EC2 images are now SP6, and the systemd-sysvcompat package is installed on SLES, so that local boot-time scripts continue to work.                                                                                                                                                                           |           |

## Changes

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Addresses |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Use the latest Barman from PGDG on RHEL-like systems. TPA previously defaulted to Barman 3.9 when installing from PGDG on a RHEL-like system, as a workaround for broken packages. More recent barman packages are OK, so we now let yum install the latest packages.                                                                                                                                                                                                                                                                                                                                                                                  |           |
| Set up EDB repositories via setup script for RedHat and Debian. This change simplifies EDB repository setup on TPA nodes into a single task for <code>dnf</code> , <code>yum</code> and <code>apt</code> package managers. This should reduce the occurrence of unexpected HTTP errors when adding EDB repositories to target hosts.                                                                                                                                                                                                                                                                                                                   |           |
| Fixed documentation for <code>efm_conf_settings</code> . Previously, documentation stated <i>"You can use <code>efm_conf_settings</code> to set any parameters, whether recognized by TPA or not. Where needed, you need to quote the value exactly as it would appear in <code>efm.properties</code>"</i> However, the <code>efm.properties.j2</code> template uses the values from <code>efm_conf_settings</code> as an Ansible dictionary, so the entries must be written in <code>key: value</code> form. <code>yaml cluster_vars: efm_conf_settings: notification.level: WARNING ping.server.ip: &lt;well known address in the network&gt;</code> |           |
| Bump Python dependency to version 3.12. TPA now requires Python 3.12. When installing TPA from EDB Repos 2.0, an appropriate Python interpreter will be installed automatically.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |           |
| Wait for protocol version update during PGD upgrade. During upgrade from PGD3 to PGD5, the protocol version update may take some time. The PGD 5 specific config changes will fail if they are attempted before the protocol version change. We have added a wait to avoid such failures.                                                                                                                                                                                                                                                                                                                                                              |           |

continues on next page

Table 20 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Addresses |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Update <code>&lt;clustername&gt;.nodes</code> when new nodes are added to an existing EFM cluster. In previous versions of TPA, when a new EFM node was added to <code>config.yml</code> , it was not listed in the <i>Allowed node host list</i> on the existing EFM nodes in the cluster. The task which executes <code>efm upgrade-conf</code> and propagates the changes from <code>/raw/&lt;clustername&gt;.properties</code> and <code>/raw/&lt;clustername&gt;.nodes</code> is now run when EITHER of these files are changed. This results in the new EFM node being written to the <code>&lt;clustername&gt;.nodes</code> file and <code>efm cluster-status</code> including it in the <i>Allowed node host list</i> . |           |
| Include <code>PGPORT</code> in the postgres user's <code>.bashrc</code> file. The <code>PGPORT</code> environment variable will now be exported as part of the postgres user's <code>.bashrc</code> file. It defaults to the port value used by the selected <code>postgres_flavour</code> , or <code>postgres_port</code> if specified in the config file.                                                                                                                                                                                                                                                                                                                                                                     |           |

## Bug Fixes

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Addresses |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Fixed an issue whereby databases were created with the default parameters rather than as configured. In previous versions of TPA, in order to suppress the <code>CREATE EXTENSION</code> statement for extensions that do not require it (aka 'modules'), the entire <code>postgres_databases</code> hash was modified. This introduced a bug, since a new hash was created that ONLY contained the database <i>name</i> and list of <i>extensions</i> , ignoring all other configuration settings for the database ( <i>owner</i> , <i>template</i> , <i>encoding</i> etc). This resulted in databases being created with the default parameters rather than as configured. To fix this, the modules are removed from the list of extensions and the resulting list is passed to the task which runs <code>CREATE EXTENSION</code> .                | 44539     |
| Fixed an issue with <code>shared_preload_libraries</code> on Patroni clusters. Entries in <code>shared_preload_libraries</code> are now treated correctly by Patroni. This fixes a bug whereby adding the pglogical extension to a Patroni cluster via <code>config.yml</code> would fail.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |           |
| Fixed a bug whereby deploys would fail for distributions which have no additional repository setup commands for extensions. In previous versions of TPA, when <code>postgres</code> was added to <code>extra_postgres_extensions</code> or the <code>extensions</code> list of a database in <code>postgres_databases</code> , deploys would fail for Debian, SLES and Ubuntu because their list of <code>repository_setup_commands</code> was empty (only RHEL has an additional command to run <code>'crb enable'</code> ). This empty list was passed to the <code>command</code> module, which would fail with <i>no command given</i> , resulting in deployment failure. The <i>Automatically run additional repository setup commands for recognized extensions</i> task is now skipped if the distribution has no additional commands to run. |           |
| Fixed an issue whereby TPA attempted to create replication slots even when <code>rep_mgr_use_slot</code> was set to 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 41776     |

continues on next page

Table 21 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Addresses |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Set <i>bdr_client_dsn_attributes</i> as the default for <i>pgd_proxy_dsn_attributes</i> and <i>pgd_cli_dsn_attributes</i> . Because pgd-proxy and pgd-cli are written in Go and use a Go driver, they do not support the full set of parameter keywords supported by libpq. In the case a cluster has installed pgd-proxy and/or pgd-cli and has configured <i>bdr_client_dsn_attributes</i> with parameters that the Go driver does <i>not</i> support, two new configuration variables must be included: <i>pgd_proxy_dsn_attributes</i> and <i>pgd_cli_dsn_attributes</i> , containing only additional DSN parameters that the Go driver supports. Conversely, if pgd-proxy and pgd-cli are installed and <i>bdr_client_dsn_attributes</i> does not include any Go-incompatible parameters, the connection strings for these tools will be configured with the attributes in <i>bdr_client_dsn_attributes</i> . This amends unexpected behavior where the <i>pgd_proxy_dsn_attributes</i> and <i>pgd_cli_dsn_attributes</i> were defaulting to empty strings when not defined, even though the user was expecting the parameters in <i>bdr_client_dsn_attributes</i> to be used. | 44819     |
| Use standard form of home directory for etcd. When creating the etcd user, TPA now refers to its home directory without a trailing slash, matching the usage of other tools like 'useradd'.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |           |
| Fixed an issue whereby the <i>cluster_ssh_port</i> setting was not respected by Barman. The Barman configuration is now able to use custom ssh port set via the <i>cluster_ssh_port</i> in <i>config.yml</i> , which defaults to 22 if it is not set. The <i>-p/--port</i> flags are now included in the <i>ssh</i> command in <i>barman.d.conf</i> and <i>barman-wal-restore</i> / <i>barman-wal-archive</i> commands respectively.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |           |

## 1.2.9 Trusted Postgres Architect 23.35.0 release notes

Released: 25 November 2024

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.35.0 include the following:

### Highlights

- Options for STIG/CIS compliance.
- Support for PGD Lightweight architecture
- Postgis is now a recognized extension.
- Docker `configure` creates named networks with static IP addresses.
- Support for RedHat Enterprise Linux 9 for ARM architectures.
- Support for PostgreSQL, EDB Postgres Extended, and EDB Postgres Advanced Server 17.

### Enhancements

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Addresses |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <p>Support STIG/CIS compliance. TPA now supports command-line options to create a cluster configured to conform to many of the requirements of the STIG and CIS security standards. These options cause TPA to set postgresql.conf settings as defined in the relevant standards, to install required extensions, to configure other aspects of system behaviour such as filesystem permissions and user connection limits, and to check for other requirements such as FIPS crypto standards which TPA can't directly impose. The clusters thus generated are not certified by TPA to conform to the standards, but much of the ground-work of creating a conforming cluster is now automated.</p>                                                                                                                                                                                                         |           |
| <p>Add support for PGD Lightweight architecture. TPA is now able to generate a PGD Lightweight architecture comprised of three nodes in two locations (2 nodes in Primary and one in Disaster Recovery) designed to ease migrations from physical replication. Users can now run <i>tpaexec configure lw -a Lightweight --postgresql 15</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |           |
| <p>Have <i>configure</i> create a user-defined network on Docker. The <i>configure</i> command will now automatically add a named network and static IP addresses to <i>config.yml</i> when Docker is the selected platform. The network name is the same as the cluster name and the address range follows the existing semantics of the <i>--network</i> option with the exception that only one subnet is used for the whole cluster rather than one per location. If a subnet prefix is not specified by the user, TPA will attempt to select a prefix which results in a subnet large enough to fit the whole cluster. The key <i>ip_address</i> may now be used to specify a static IP for a Docker instance as long as a named network is specified in the <i>config.yml</i>.</p>                                                                                                                    |           |
| <p>Added experimental support for using an existing Barman node as backup node in new cluster. When using an existing Barman node as a backup node in a new cluster, users can set <i>barman_shared: true</i> in the Barman instance's vars with the platform set to <i>bare</i> and other information supplied as usual for bare instances. This change allows TPA to skip some configuration steps that would otherwise fail due to usermod issues, as the Barman user already has running processes from previous deployments. The shared Barman instance is treated as a bare instance, so the required access, including the Barman user's access to the target PostgreSQL instances, must be already in place. Copying the Barman user's keys from the original cluster to the new cluster can be used to achieve this, see the Barman section of the TPA documentation for detailed information.</p> |           |
| <p>Add <i>postgis</i> to list of recognized extensions. The PostGIS package will automatically be added when a user specifies <i>postgis</i> as an entry in either <i>postgres_extensions</i> or the list of extensions named under <i>postgres_databases</i>. Also enables the CRB (Code Ready Builder) repository for RHEL-compatible distributions so PostGIS dependencies can be installed.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |           |
| <p>Enable EFM probes when a PEM agent is registered on an EFM node. The <i>--efm-install-path</i> and <i>--efm-cluster-name</i> flags are set when a PEM server is registered on an EFM node. The <i>Streaming Replication</i>, <i>Failover Manager Node Status</i> and <i>Failover Manager Cluster Info</i> probes are enabled when a PEM agent is registered on an EFM node.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |           |
| <p>Support RedHat Enterprise Linux 9 for ARM architectures. Packages are now published targeting RHEL 9 ARM64, and TPA supports deployments using this architecture and OS. Also updated the list of supported AWS images to include the RedHat 9 ARM64 AMI provided by Amazon. The default <i>instance_type</i> for ARM64 EC2 instances has been updated from <i>a1</i> to <i>t4g</i>, which is the current generation processor available for burstable general purpose workloads.</p>                                                                                                                                                                                                                                                                                                                                                                                                                    |           |

continues on next page

Table 22 – continued from previous page

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Addresses |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Support PostgreSQL, EDB Postgres Extended, and EDB Postgres Advanced Server 17. Clusters can be configured to use PostgreSQL, EDB Postgres Extended and EDB Postgres Advanced Server version 17. Barman no longer needs to install the postgres server package to get the <i>pg_receivewal</i> binary when using EDB Postgres Advanced Server 17 or EDB Postgres Extended 17 since the binary has been added to the client package for these versions. TPA raises an architecture error when a cluster is configured with <i>repmgr</i> as the failover_manager as it is not available for Postgres 17. Updated documentation to reflect supported versions.      |           |
| Make <i>password_encryption</i> algorithm for <i>efm</i> Postgres user configurable. Expose a configurable <i>efm_user_password_encryption</i> variable which should be set to either 'md5' or 'scram-sha-256' depending on user requirements. This controls the <i>auth-method</i> for the <i>efm</i> Postgres user in <i>pg_hba.conf</i> and the algorithm used for generating it's encrypted password. In clusters deployed with <i>compliance</i> configured to <i>stig</i> , the 'efm' Postgres user's <i>auth-method</i> in <i>pg_hba.conf</i> will be set to <i>scram-sha-256</i> since FIPS-enabled operating systems do not allow <i>md5</i> to be used. |           |
| Allow multiple addresses to be supplied with hostnames. When using the <i>--hostnames-from</i> option to <i>tpaexec configure</i> , you can now include two ip addresses on each line, which will be included in the generated config.yml as <i>public_ip</i> and <i>private_ip</i> .                                                                                                                                                                                                                                                                                                                                                                             |           |

## Changes

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Addresses |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Remove deprecated <i>PermissionStartOnly</i> in postgres.service.j2 template. <i>PermissionStartOnly</i> has been deprecated and is now achieved via <i>ExecStartPost=+/bin/bash...</i> syntax                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |           |
| The <i>barman</i> Postgres user is no longer a superuser. Certain required privileges are granted to Postgres role, <i>barman_role</i> , which is then granted to the <i>barman</i> Postgres user. This avoids creating the <i>barman</i> user as a superuser. This role can also be granted to other Postgres users by adding it to their <i>granted_roles</i> list using <i>postgres/createuser</i> . The <i>barman_role</i> is created as part of the Barman tasks; if Barman is not used, this role will not be created. Therefore, the task that grants privileges to this role is only executed if the <i>barman_role</i> username is in the list of Postgres users that are created. The 'barman' user now has <i>NOSUPERUSER</i> explicitly specified as a role attribute. If a cluster was deployed with a previous TPA version (which created the 'barman' user as a superuser), deploying with this version will remove the <i>superuser</i> role attribute from the <i>barman</i> user. |           |
| Add new option <i>harp_local_etcd_only</i> when using etcd with HARP. Add new optional var <i>harp_local_etcd_only</i> available when using etcd with HARP. This option tells HARP manager to connect to local etcd node. This recommendation follows the best practices learnt by doing the same when <i>bdr</i> as consensus protocol is being used. The default mode of adding multiple endpoints can lead to performance issues in some cases. This option is added to give more control to the user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |           |
| Improve postgres-monitor script. Improve postgres-monitor script to better manage recoverable errors and add retries on network errors to ensure that it won't return failure when it just didn't allow enough time for postgres service to be fully started.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |           |
| Only add nodes with <i>efm</i> role to cluster <i>efm.nodes</i> file. Previously the <i>pemserver</i> and <i>barman</i> nodes were added to the <i>Allowed node host list</i> in EFM when they were not relevant to EFM functions. Refactored the task that writes the <i>efm.node</i> configuration to only include those nodes that have <i>efm</i> in their list of roles.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |           |

## Bug Fixes

| Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Addresses |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Fix tpaexec test for pgd-proxy config verification. Fixed a bug whereby the test that ensures the current pgd-proxy configuration matches the expected configuration would fail for version < 5.5.0. This fix ensures that TPA won't try to query configuration keys added in version 5.5.0.                                                                                                                                                                                                                                                                          |           |
| Fix case where <i>primary_slot_name</i> added for EFM compatibility interferes with <i>bdr_init_physical</i> . A <i>primary_slot_name</i> is configured on the primary node to ensure the old primary uses a physical slot for replication during an EFM switchover. However, 'bdr_init_physical' attempts to use it for node initialisation and hangs indefinitely since the slot does not exist in a PGD installation. This <i>primary_slot_name</i> is now conditionally set explicitly when the <i>failover_manager</i> is EFM to avoid setting it unnecessarily. |           |
| Download correct <i>bash-completion</i> package version. If the <i>pgdcli_package_version</i> is specified in <i>config.yml</i> , the <i>bash-completion</i> package is incorrectly named because the <i>packages_for</i> filter erroneously appends the <i>pgdcli_package_version</i> to the package name. This results in an attempt to download a nonexistant package. The <i>bash-completion</i> package is now appended to the list after the <i>packages_for</i> filter, since it's version is independent from the <i>pgdcli_package_version</i> .             |           |
| Fix an issue whereby in some cases error messages would be repeated even after successful tasks. TPA now clears the error message stack after each task to ensure messages are not spuriously repeated                                                                                                                                                                                                                                                                                                                                                                |           |
| Fix issue that prevented the addition of replicas to Patroni clusters. Fixed an issue whereby new replicas in Patroni clusters would fail with errors related to replication slots.                                                                                                                                                                                                                                                                                                                                                                                   |           |
| Add <i>pem-agent</i> role on barman nodes at most once for M1 architecture. If <i>--enable-pem</i> and <i>--enable-pg-backup-api</i> are passed to <i>tpaexec configure</i> , <i>pem-agent</i> is added twice to the <i>barman</i> node if it is also a <i>witness</i> . Fixed by consolidating both <i>if</i> statements together to only evaluate the conditions once.                                                                                                                                                                                              |           |
| Set <i>pem_python_executable</i> outside of the <i>pkg</i> role. Fixed a bug whereby if the user excluded the <i>pkg</i> selector, later PEM-related tasks would fail because the <i>pem_python_executable</i> fact had not been set.                                                                                                                                                                                                                                                                                                                                 |           |

## 1.2.10 Trusted Postgres Architect 23.34.1 release notes

Released: 9 September 2024

Trusted Postgres Architect 23.34.1 is a bug fix release which resolves the following issues:

| Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bug Fix | Fixed an issue whereby running deploy after a switchover fails for nodes with <i>efm-witness</i> role. The <i>upstream-primary</i> for EFM nodes is determined using the facts gathered from Postgres. This previously failed for nodes with <i>efm-witness</i> roles since they do not have Postgres. The task to determine upstream-primary is now run only on nodes with <i>primary</i> or <i>replica</i> roles. |

### 1.2.11 Trusted Postgres Architect 23.34 release notes

Released: 22 August 2024

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.34 include the following:

| Type        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enhancement | TPA now has an <i>efm-pre-config</i> hook which runs after efm has been installed and its configuration directory and user have been created, but before efm is configured. It can be used to install custom efm helper scripts.                                                                                                                                                                                                                                        |
| Enhancement | TPA now has its own output plugin, which shows one line of information per task, omitting tasks for which even one line would be uninformative. The lines are indented to enable TPA's control flow to be visible, and include color-coded counts of successful, skipped, and ignored hosts. The fuller output can be turned on by setting <code>TPA_USE_DEFAULT_OUTPUT=true</code> in your environment, or by adding the <code>-v</code> switch to the command line.   |
| Enhancement | TPA now allows you to specify additional options when registering PEM agents by listing them under <i>pemagent_registration_opts</i> in <i>cluster_vars</i> .                                                                                                                                                                                                                                                                                                           |
| Enhancement | You can now provide your own web server certificates for use with the PEM server by including the names of the certificate and key pair for use on the PEM server in <i>config.yml</i> under the <i>cluster_vars</i> or <i>pem-server</i> instance vars <i>pem_server_ssl_certificate</i> and <i>pem_server_ssl_key</i> . TPA will copy them from the <i>ssl/pemserver</i> directory of the cluster directory to the PEM server and configure Apache/httpd accordingly. |
| Enhancement | TPA now runs the EFM <i>upgrade-conf</i> command on new cluster deployments to benefit from the comments and inline documentation that are added to both <i>&lt;cluster_name&gt;.properties</i> and <i>&lt;cluster_name&gt;.nodes</i> files.                                                                                                                                                                                                                            |
| Enhancement | TPA can now set up EFM clusters using hostname resolution instead of IP addresses for <i>bind.address</i> value. This can be invoked with the <i>--efm-bind-by-hostname</i> option for the configure command or <i>efm_bind_by_hostname: true/false</i> in <i>cluster_vars</i> section of <i>config.yml</i> .                                                                                                                                                           |
| Enhancement | TPA now supports setting the EFM properties that added in EFM 4.9: <i>enable.stop.cluster: boolean, default true, priority.standbys: default "</i> , <i>detach.on.agent.failure: boolean, default false, pid.dir: default "</i> .                                                                                                                                                                                                                                       |
| Enhancement | TPA can now configure pgBouncer to use <i>cert</i> authentication for connections from pgBouncer to Postgres. This is enabled by setting <i>pgbouncer_use_cert_authentication</i> to true in <i>cluster_vars</i> . When enabled, the authentication method for users connecting to pgBouncer is also changed from <i>md5</i> to <i>scram-sha-256</i> .                                                                                                                  |
| Change      | TPA no longer supports RAID creation on AWS.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Change      | Removed EFM dependency for resolving <i>upstream_primary</i> . Previously, EFM was queried for the current primary on a deploy after a switchover. If EFM is not running, this will fail. Now the <i>cluster_facts</i> collected through Postgres are used to determine the current primary after a switchover, removing the dependency on EFM.                                                                                                                         |
| Change      | In EFM clusters, the <i>upstream_primary</i> is now correctly updated after switchover, resulting in the correct <i>auto.reconfigure</i> setting be set on replicas. Standbys now follow the new primary.                                                                                                                                                                                                                                                               |
| Bug Fix     | Fixed an issue whereby TPA would incorrectly apply proxy settings when accessing the Patroni API. The Ansible default is to use a proxy, if defined. This does not work in the (rather common) case of an airgapped environment that needs a proxy to download packages from the internet, because the proxy also intercepts (and disrupts) calls to the Patroni API.                                                                                                   |

continues on next page

Table 26 – continued from previous page

| Type    | Description                                                                                                                                                                                                                                           |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bug Fix | Fixed an issue whereby TPA would set PEM agent parameters on all instances that were only appropriate for the <i>pemserver</i> instance.                                                                                                              |
| Bug Fix | Added missing entries for pgd-proxy and pgdcli default package name when using SLES operating system as target for cluster nodes.                                                                                                                     |
| Bug Fix | Fix an issue whereby TPA would fail to reload/restart postgres on existing nodes to re-read configuration changes and the new node would therefore fail to connect to the cluster.                                                                    |
| Bug Fix | Fixed an issue whereby when taking backups from a replica, barman could fail when taking its initial backup by timing out waiting for WAL files. This is fixed by waiting for barman to complete its base backup before forcing a WAL segment switch. |
| Bug Fix | Ensure that <i>repmgr witness register</i> command is used with the correct postgres_port value even when using non-default postgres port for the upstream_primary postgres.                                                                          |
| Bug Fix | Fixed an issue whereby failover_manager override to <i>repmgr</i> would not work correctly when set at instance level for subscriber-only nodes and their replicas in PGD clusters.                                                                   |
| Bug Fix | Fixed two cases of incorrect cgroup detection: on MacOSX, we no longer try to read <i>/proc/mounts</i> . On systems where <i>/sys/fs/cgroup</i> is <i>ro</i> but mounts under it are <i>rw</i> , TPA now correctly detects this.                      |
| Bug Fix | Ensure we can verify the actual config set on pgd-proxy nodes for the newly added <i>read_listen_port</i> option in pgd-proxy.                                                                                                                        |
| Bug Fix | Fixed an issue that would prevent deployment with PEM 9.7.0. PEM 9.7.0 no longer depends on Apache at a package level therefore to use Apache as the web server we install the packages explicitly.                                                   |

## 1.2.12 Trusted Postgres Architect 23.33 release notes

Released: 24 Jun 2024

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.33 include the following:

| Type        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enhancement | TPA now supports Debian 12 Bookworm on the ARM64 CPU architecture.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Enhancement | PGD version 5.5 allows for proxy nodes to be configured as read endpoints, which direct read-only queries to a shadow node. TPA supports this configuration option by setting a <i>read_listen_port</i> parameter under <i>default_pgd_proxy_options</i> and <i>pgd_proxy_options</i> in <i>config.yml</i> . This parameter is included by default when the PGD version is 5.5 or greater. Users can also specify the port numbers by passing <i>--proxy-listen-port</i> and <i>proxy-read-listen-port</i> arguments to the <i>tpaexec configure</i> command. |
| Enhancement | TPA now supports deployment and configuration of the Beacon Agent on any Postgres node by assigning the role 'beacon-agent' or using the <i>--enable-beacon-agent</i> option with <i>configure</i> .                                                                                                                                                                                                                                                                                                                                                          |
| Enhancement | Added support for <i>postgres_wal_dir</i> in Patroni deployments. When a custom <i>postgres_wal_dir</i> is specified in TPA configuration, TPA will make sure to relay that option to the corresponding settings in the Patroni configuration file. That way, if Patroni ever needs to rebuild a standby on its own, out of TPA, the standby will be properly set up with a custom WAL directory.                                                                                                                                                             |
| Enhancement | When adding PgBouncer nodes in a Patroni cluster, TPA now configures Patroni with a <i>on_role_change</i> callback. That callback takes care of updating the primary connection info in the PgBouncer nodes in response to failover and switchover events.                                                                                                                                                                                                                                                                                                    |

continues on next page

Table 27 – continued from previous page

| Type        | Description                                                                                                                                                                                                                                                                                                                                                         |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enhancement | EDB now produces its own <i>edb-patroni</i> package instead of rebuilding the <i>patroni</i> packages from PGDG. TPA now allows users to select between <i>patroni</i> and <i>edb-patroni</i> packages. The selection is made through the new TPA setting <i>patroni_package_flavour</i> .                                                                          |
| Change      | To work around broken Barman 3.10 packages in the PGDG repos, TPA now installs version 3.9 of Barman if using PGDG repos on an RHEL-family system. This behavior can be overridden by explicitly setting <i>barman_package_version</i> in <i>config.yml</i> .                                                                                                       |
| Change      | The <i>haproxy_bind_address</i> is now set to <i>0.0.0.0</i> when Patroni is the failover manager. This resolves an issue with the general default of <i>127.0.0.1</i> preventing communication between Postgres nodes and HA Proxy nodes. Users should change this value to something more restrictive and appropriate for their cluster networking.               |
| Change      | Task selectors are now consistently applied in the final stage of deployment. Consistency of task selectors in the tests is improved and the examples of task selectors in the docs are now correct. All deploy-time hooks now have corresponding task selectors.                                                                                                   |
| Change      | If <i>barman_package_version</i> is set, TPA will now look at it when looking for the <i>barman-cli</i> package as well as for Barman itself. This resolves an inconsistency which caused clusters using the downloader to fail when <i>barman_package_version</i> was used.                                                                                        |
| Bug Fix     | Fixed an issue whereby required permissions on functions in the BDR database were not being granted to the HARP DCS user on a witness node.                                                                                                                                                                                                                         |
| Bug Fix     | Fixed an issue whereby docker provisioning failed with "read-only file system". On host systems running <i>cgroup1</i> with docker containers running recent OS images, <i>tpaexec provision</i> could fail to provision containers with an error message like "mkdir /sys/fs/cgroup/tpa.scope: read-only file system". TPA will now detect this case and avoid it. |
| Bug Fix     | TPA now provides a clear error message if the user runs <i>tpaexec cmd</i> or <i>tpaexec ping</i> before provision.                                                                                                                                                                                                                                                 |

### 1.2.13 Trusted Postgres Architect 23.32 release notes

Released: 15 May 2024

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.32 include the following:

| Type        | Description                                                                                                                                                                                                                                                                                                                                                                    |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enhancement | The M1 architecture now supports the following additional arguments to <i>tpaexec configure</i> : <i>--location-names</i> , <i>--primary-location</i> , <i>--data-nodes-per-location</i> , <i>--witness-only-location</i> , and <i>--single-node-location</i> . By combining these arguments, most common layouts can be specified without needing to edit <i>config.yml</i> . |
| Enhancement | TPA now installs chrony during deploy, keeping the default config upon all except on AWS where we point to Amazon Time Sync service.                                                                                                                                                                                                                                           |
| Enhancement | TPA now supports RHEL 8 and 9 on IBM Power (PPC64le).                                                                                                                                                                                                                                                                                                                          |
| Enhancement | Added a <i>--force</i> option to <i>tpaexec relink</i> . By default, relink doesn't modify targeted files if they already exist. With <i>--force</i> , relink removes all existing targeted files then recreates them.                                                                                                                                                         |
| Enhancement | TPA now supports Debian 12 x86.                                                                                                                                                                                                                                                                                                                                                |
| Enhancement | <i>pg_failover_slots</i> is now a recognized extension                                                                                                                                                                                                                                                                                                                         |
| Enhancement | The <i>sql_profiler</i> , <i>edb_wait_states</i> and <i>query_advisor</i> extensions are now automatically included for any <i>pem-agent</i> node. The list of default extensions for <i>pem-agent</i> nodes can be overridden by including a list of <i>pemagent_extensions</i> in <i>config.yml</i> . If this list is empty, no extensions will be automatically included.   |

continues on next page

Table 28 – continued from previous page

| Type          | Description                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Change        | TPA can now provision Docker clusters on hosts running cgroups 2 for all systems except RHEL 7. On newer systems (RHEL 9 or Ubuntu 22.04), TPA will use cgroups 2 scopes for additional isolation between the host and the containers.                                                                                                                                                                                                   |
| Change        | Updated AWS AMI versions to the latest versions.                                                                                                                                                                                                                                                                                                                                                                                         |
| Bug Fix       | Fixed an issue whereby deploying to Debian 10 on AWS would fail with the message <i>The repository 'http://cdn-aws.deb.debian.org/debian buster-backports Release' does not have a Release file.</i> The backports repository for debian 10 (buster) is no longer available on deb.debian.org but the standard AWS AMI still refers to it, so we modify <code>/etc/apt/sources.list</code> accordingly before attempting apt operations. |
| Bug Fix       | Fixed an issue whereby deployment would fail on AWS when <code>assign_public_ip:no</code> was set.                                                                                                                                                                                                                                                                                                                                       |
| Bug Fix       | Fixed problems with various roles that caused mixed errors when trying to use custom users for barman and postgres, thereby resulting in a failed deployment.                                                                                                                                                                                                                                                                            |
| Bug Fix       | Fixed an issue whereby deployments after the initial one could fail with an error like <i>Unrecognised host=... in primary_conninfo</i> if the key <code>ip_address</code> was used to define the IP address.                                                                                                                                                                                                                            |
| Bug Fix       | Fixed an error whereby <code>tpaexec upgrade</code> could invoke the relink script in a way which caused an error and showed an unhelpful usage message for <code>tpaexec relink</code> .                                                                                                                                                                                                                                                |
| Bug Fix       | Fixed an issue whereby a task to reload Postgres was skipped resulting in the <code>restore_command</code> override not being removed from <code>postgresql.auto.conf</code>                                                                                                                                                                                                                                                             |
| Bug Fix       | Fixed an issue whereby TPA did not change to the source directory before attempting to compile BDR from source.                                                                                                                                                                                                                                                                                                                          |
| Bug Fix       | Fixed an issue whereby TPA would require a valid 2ndQuadrant token even if one was not needed for the specified cluster.                                                                                                                                                                                                                                                                                                                 |
| Documentation | Updated the tower/AAP documentation to include instructions on creating an AAP Execution Environment.                                                                                                                                                                                                                                                                                                                                    |

### 1.2.14 Trusted Postgres Architect 23.31 release notes

Released: 19 Mar 2024

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.31 include the following:

| Type    | Description                                                                |
|---------|----------------------------------------------------------------------------|
| Bug Fix | Fixed a critical bug whereby deployments could fail due to a syntax error. |

### 1.2.15 Trusted Postgres Architect 23.30 release notes

Released: 19 Mar 2024

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.30 include the following:

| Type        | Description                                                                                                                                                                                                                                                                             |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| New Feature | TPA now provides a custom 'Execution Environment' image to be used in Ansible Automation Platform 2.4+ (Controller version 4+). this image contains everything needed to run deployments via AAP. This image is built using ansible-builder and a python-alpine lightweight base image. |

continues on next page

Table 30 – continued from previous page

| Type        | Description                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enhancement | TPA now automatically adds package names and shared preload library entries for a subset of extensions. For these specific extensions, only the extension name is needed in the <i>extra_postgres_extensions</i> list or the <i>extensions</i> list of a database entry in <i>postgres_databases</i> .                                                                                                                |
| Enhancement | The EDB Advanced Storage Pack package and shared preload library entry will automatically be added for <i>bluefin</i> when a user specifies it as an extension and the <i>postgres_version</i> is 15 or greater.                                                                                                                                                                                                      |
| Enhancement | Added a new 'provision_only' option for instances. If an instance has <i>provision_only: true</i> in <i>config.yml</i> , it will be provisioned as normal but not added to the inventory which is seen by <i>tpaexec deploy</i> .                                                                                                                                                                                     |
| Change      | Previous versions of TPA used to synchronize the source node's database structure to witness nodes. This was not necessary and the synchronized schema was never be used or updated. To prevent this happening, TPA now explicitly sets "synchronize_structure" to "none" when calling <i>bdr.join_node_group()</i> for witness nodes.                                                                                |
| Change      | Selective execution of tasks is now supported using custom selectors rather than Ansible tags. To run only tasks matching a certain selector: <i>tpaexec deploy --included_tasks=barman</i> . To skip tasks matching a certain selector: <i>tpaexec deploy --excluded_tasks=ssh</i> Task selectors can also be used by specifying the <i>excluded_tasks</i> or <i>included_tasks</i> variables in <i>config.yml</i> . |
| Change      | Ansible 2.9 is no longer supported, neither the community distribution nor the 2ndQuadrant fork. Users who have been using the <i>--skip-tags</i> option to <i>tpaexec deploy</i> should move to the new <i>--excluded_tasks</i> option.                                                                                                                                                                              |

## 1.2.16 Trusted Postgres Architect 23.29 release notes

Released: 15 Feb 2024

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.29 include the following:

| Type        | Description                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enhancement | Added support for storing the cluster vault password in the system keyring. This leverages python keyring module to store vault password in the supported system keyring when <i>keyring_backend</i> is set to <i>system</i> (default for new clusters). This change does not impact existing clusters or any clusters that set <i>keyring_backend</i> to <i>legacy</i> in <i>config.yml</i> . |
| Enhancement | The <i>--ansible-version</i> argument to <i>tpaexec setup</i> now accepts 8 or 9 as valid ansible versions, as well as the existing <i>2q</i> or <i>community</i> , both of which imply ansible 2.9. The default is now 8. Support for ansible 9 is experimental and requires python 3.10 or above.                                                                                            |
| Bug Fix     | Fixed an issue whereby <i>edb_repositories</i> already defined in <i>config.yml</i> are not kept during reconfigure. Fixes <i>bdr4</i> to <i>pgd5</i> upgrade scenario in air gapped environment.                                                                                                                                                                                              |
| Bug Fix     | TPA's <i>postgres-monitor</i> will now recognize the message "the database system is not yet accepting connections" as a recoverable error.                                                                                                                                                                                                                                                    |
| Bug Fix     | TPA now correctly skips the <i>postgres/config/final</i> role on replicas when upgrading.                                                                                                                                                                                                                                                                                                      |
| Bug Fix     | Fixed an issue whereby wildcards in package names were not respected when using package downloader on Debian and Ubuntu systems.                                                                                                                                                                                                                                                               |
| Bug Fix     | The downloader now runs <i>apt-get update</i> before fetching packages on Debian and Ubuntu systems.                                                                                                                                                                                                                                                                                           |
| Bug Fix     | TPA now disables transaction streaming when CAMO is enabled in PGD clusters.                                                                                                                                                                                                                                                                                                                   |

continues on next page

Table 31 – continued from previous page

| Type    | Description                                                           |
|---------|-----------------------------------------------------------------------|
| Bug Fix | TPA now correctly configures Barman servers where selinux is enabled. |

## 1.2.17 Trusted Postgres Architect 23.28 release notes

Released: 23 Jan 2024

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.28 include the following:

| Type        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enhancement | Added a new option <i>postgres_log_file</i> . This option sets the Postgres log file, whether logging through stderr or syslog. The default is '/var/log/postgres/postgres.log', the previously hard-coded value.                                                                                                                                                                                                                                                                                           |
| Enhancement | Added a new hook <i>barman-pre-config</i> . This hook is invoked after Barman is installed and its user is set up but before it is configured. It can be used for installing certificate files or other tasks which need the barman user to exist but which must be done before Barman is started.                                                                                                                                                                                                          |
| Enhancement | The key <i>elastic_ip</i> on an AWS instance in <i>config.yml</i> can be set to an elastic IP address that has already been allocated in order to assign it to this instance.                                                                                                                                                                                                                                                                                                                               |
| Change      | In Patroni clusters, TPA now sets up replicas before handing over control of the cluster to Patroni, rather than setting up the primary only and letting Patroni set up the replicas.                                                                                                                                                                                                                                                                                                                       |
| Change      | For new clusters, TPA will create the user specified by setting <i>harp_manager_user</i> (by default <i>harpmanager</i> ), belonging to the <i>bdr_superuser</i> role, and set HARP Manager to operate as this user instead of <i>postgres</i> superuser. This does not affect the existing clusters where TPA will keep using <i>postgres</i> as the HARP Manager user, unless the user overrides this behavior by explicitly setting <i>harp_manager_user</i> to a different value in <i>config.yml</i> . |
| Bug Fix     | Fixed an issue whereby TPA would erroneously attempt to install repmgr on an EFM cluster.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Bug Fix     | Fixed an issue whereby the TPA would return a non-zero exit code when the warning about 2q repositories was displayed despite deploy having succeeded.                                                                                                                                                                                                                                                                                                                                                      |
| Bug Fix     | TPA will now interpret wildcards correctly on Debian-family systems when downloading packages for offline use.                                                                                                                                                                                                                                                                                                                                                                                              |
| Bug Fix     | Fixed an issue whereby TPA would attempt to use incorrect package names for repmgr when installing from PGDG repositories.                                                                                                                                                                                                                                                                                                                                                                                  |
| Bug Fix     | Fixed barman connection failure when using selinux and a custom barman home directory.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Bug Fix     | TPA will now use the correct cluster name in <i>show-password</i> and <i>store-password</i> commands when it is different from the directory name                                                                                                                                                                                                                                                                                                                                                           |
| Bug Fix     | TPA will now error out cleanly if unavailable 2ndQuadrant repository keys are required.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Bug Fix     | TPA will now sanitize hostnames correctly when the <i>--cluster-prefixed-hostnames</i> option is used.                                                                                                                                                                                                                                                                                                                                                                                                      |
| Bug Fix     | TPA will now ensure packages are correctly copied to the remote host when upgrading a cluster using a local repo.                                                                                                                                                                                                                                                                                                                                                                                           |

### 1.2.18 Trusted Postgres Architect 23.27 release notes

Released: 19 Dec 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.27 include the following:

| Type        | Description                                                                                                                                                                                                                                                 |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enhancement | TPA now supports Oracle Linux 7, 8 and 9 on Docker.                                                                                                                                                                                                         |
| Change      | TPA now requires Python 3.9-3.11 and depends on the package <i>edb-python-39</i> to provide a suitable interpreter.                                                                                                                                         |
| Change      | TPA will no longer configure any 2ndQuadrant repositories by default, instead it will select suitable repositories from EDB Repos 2.0.                                                                                                                      |
| Change      | TPA now provides a new <i>--replace-2q-repositories</i> argument to <i>tpaexec reconfigure</i> that will remove 2q repositories from an existing config.yml and add suitable EDB repositories for the cluster's postgres flavour and BDR version.           |
| Change      | TPA now sets file system permissions explicitly on more objects.                                                                                                                                                                                            |
| Change      | A new variable <i>disable_repository_checks</i> can be set to true in config.yml to bypass the usual check for EDB repositories when deploying the PGD-Always-ON architecture.                                                                              |
| Change      | TPA will now generate a <i>primary_slot_name</i> also on primary node to be used in case of switchover, to ensure the switched primary will have a physical slot on the new primary.                                                                        |
| Change      | TPA will now ensure that <i>commit_scope</i> for CAMO enabled partners is generated using existing config options from older BDR versions when running <i>tpaexec reconfigure</i> command to prepare for major PGD upgrade. It also choses better defaults. |
| Bug fix     | Fixed an issue whereby postgres variables were rejected by Patroni due to validation rules.                                                                                                                                                                 |
| Bug fix     | Fixed an issue whereby a user could not set a single <i>barman_client_dsn_attributes</i> with <i>sslmode=verify-full</i> .                                                                                                                                  |
| Bug Fix     | TPA will now assign a lower default <i>maintenance_work_mem</i> to avoid out-of-memory errors.                                                                                                                                                              |

### 1.2.19 Trusted Postgres Architect 23.26 release notes

Released: 30 Nov 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.26 include the following:

| Type        | Description                                                                                                                                                                   |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enhancement | TPA now supports Oracle Linux 9 on the Docker platform.                                                                                                                       |
| Enhancement | Added <i>--cluster-prefixed-hostnames</i> option to <i>tpaexec configure</i> . This makes it easy to avoid hostname clashes on machines hosting more than one docker cluster. |
| Change      | Added packages to enable Docker builds on Mac OS X.                                                                                                                           |
| Change      | When there are multiple PEM servers in a cluster, the agent running on a PEM server registers to its local server.                                                            |
| Change      | For PGD 5 clusters with CAMO. TPA will set timeout to 60s and <i>require_write_lead</i> to true by default.                                                                   |
| Bug Fix     | Fixed an issue whereby CAMO config was not correctly set up when upgrading a PGD 3 cluster to PGD 5. Upgrade is now fully supported for CAMO clusters.                        |
| Bug Fix     | Fixed an issue whereby <i>hostname</i> rather than <i>bdr_node_name</i> was used when fencing or unfencing a HARP node.                                                       |
| Bug Fix     | Fixed an issue whereby <i>provision</i> would be automatically run when <i>deploy</i> was invoked with options that suppress deployment.                                      |

## 1.2.20 Trusted Postgres Architect 23.25 release notes

Released: 14 Nov 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.25 include the following:

| Type        | Description                                                                                                                                                                 |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enhancement | TPA now supports automated upgrades from PGD 3.7 to PGD 5.3 or above. <i>Note, upgrading clusters with CAMO is not yet supported.</i>                                       |
| Enhancement | TPA now supports EDB Advanced Server 16 and EDB Extended Server 16.                                                                                                         |
| Change      | Various improvements to the upgrade process introduced with PGD 4 to PGD 5 upgrades have been backported to BDR-Always-ON upgrades.                                         |
| Change      | TPA now supports installing PEM on SLES.                                                                                                                                    |
| Change      | TPA now explicitly sets permissions when creating some filesystem objects. This will be extended to all filesystem objects in a future release.                             |
| Change      | TPA now adds a symlink to the pgd-cli config file for v1 so it can be run without having to specify the path via <code>-f</code> switch.                                    |
| Change      | TPA now calls the <code>alter_node_kind</code> PGD function to ensure node kind is set correctly for BDR-Always-ON clusters using BDR version 4.3 and above.                |
| Change      | Default cluster configuration from now selects SLES 15 SP5 when SLES 15 is requested (previously SP4).                                                                      |
| Bug Fix     | Fixed an issue which resulted in a checksum failure during <code>tpaexec setup</code> command for <code>tpaexec-deps</code> users.                                          |
| Bug Fix     | Fixed an issue whereby <code>pem_server_group</code> was not correctly applied when <code>pemworker</code> was invoked meaning servers were not grouped as expected in PEM. |
| Bug Fix     | Fixed an issue with the <code>sys/sysstat</code> role whereby <code>sar</code> was not scheduled to run on instances other than the Barman instance.                        |

## 1.2.21 Trusted Postgres Architect 23.24 release notes

Released: 17 Oct 2023

Support for using the 2ndQuadrant Ansible fork will be removed from TPA in April 2024 and the GitHub repository will be archived.

You should switch to Community Ansible, which is now the default. For the vast majority of users, this change will be transparent.

If you are using `--skip-tags` with 2ndQuadrant Ansible, be aware that this is not supported with TPA and Community Ansible. We plan to provide an alternative to `--skip-tags` compatible with Community Ansible before the removal of 2ndQuadrant Ansible.

| Type   | Description                                                                                                                                                                                                                                                                                                                                                                |
|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Change | <code>tpaexec setup</code> now defaults to using community ansible rather than 2ndQuadrant ansible. The option <code>--use-2q-ansible</code> can be used to force the use of 2ndQuadrant ansible, which is now deprecated and will be removed in a future release. If you are using <code>--skip-tags</code> , see <a href="#">Configuring lvm-direct for production</a> . |
| Change | When a repository has been removed from <code>edb_repositories</code> in <code>config.yml</code> , <code>tpaexec deploy</code> now removes it from the nodes.                                                                                                                                                                                                              |
| Change | TPA will now detect when <code>harp-proxy</code> and <code>harp-manager</code> are running on the same node and use a different config file for <code>harp-proxy</code> .                                                                                                                                                                                                  |
| Change | The <code>upgrade</code> command will now update local repositories on target instances.                                                                                                                                                                                                                                                                                   |

continues on next page

Table 36 – continued from previous page

| Type    | Description                                                                                                                                     |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Bug Fix | Fixed an issue whereby TPA did not respect <i>postgres_wal_dir</i> in <i>pg_basebackup</i> invocation                                           |
| Bug Fix | TPA will now accept <i>repmgr</i> as a failover manager for subscriber-only nodes in PGD clusters, allowing physical replication of such nodes. |
| Bug Fix | Fixed a typo which prevented TPA building Ubuntu 22.04 Docker images.                                                                           |
| Bug Fix | TPA will now reject unsupported combination of the BDR-Always-ON architecture, the EDB Postgres Extended flavour, and PEM at configure-time.    |

### 1.2.22 Trusted Postgres Architect 23.23 release notes

Released: 21 Sep 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.23 include the following:

| Type        | Description                                                                                                                                                                   |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enhancement | TPA now supports PostgreSQL 16. Please note, PostgreSQL 16 packages are not yet available in all supported repos, so not all configurations will work until this is the case. |
| Change      | When Postgres 16 or above is selected, TPA will not add any 2ndQuadrant repos by default. TPA will explicitly set <i>tpa_2q_repositories: []</i> in this case.                |
| Change      | EFM is now configured to use JDK 11 by default on platforms where it is available.                                                                                            |
| Change      | Where no EDB Repositories are use, TPA will not exclude any packages from PGDG (previously Barman and psycopg2 were excluded).                                                |
| Change      | Added package names for <i>etcd</i> and <i>Patroni</i> to support installation on SLES.                                                                                       |
| Bug Fix     | Fixed an issue whereby Apache HTTPD service for PEM Server would not start on boot.                                                                                           |
| Bug Fix     | Fixed an issue whereby <i>pg_backup_api</i> tests were run with incorrect permissions causing them to fail.                                                                   |
| Bug Fix     | Fixed an issue whereby Apache HTTPD service for <i>pg_backup_api</i> would not start on boot.                                                                                 |
| Bug Fix     | Fixed an issue whereby <i>bdr.standby_slot_names</i> and <i>bdr.standby_slots_min_confirmed</i> checks used the incorrect schema on bdr3 clusters.                            |
| Bug Fix     | Fixed an issue whereby configuration keys for extensions were passed to Patroni in the incorrect format, resulting in 'WARNING: Removing unexpected parameter'.               |
| Bug Fix     | Fixed an issue when using the intermediate base image option for <i>docker_images</i> whereby the resulting image name was incorrect.                                         |

### 1.2.23 Trusted Postgres Architect 23.22 release notes

Released: 6 Sep 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.22 include the following:

| Type   | Description                                                                                                       |
|--------|-------------------------------------------------------------------------------------------------------------------|
| Change | TPA is now an open source project! You can clone the source under the GPLv3 license from <a href="#">GitHub</a> . |

### 1.2.24 Trusted Postgres Architect 23.21 release notes

Released: 5 Sep 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.21 include the following:

| Type    | Description                                                                                                                                                                                                                                   |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Change  | The default M1 configuration now uses EDB Repos 2.0 if any EDB software is selected, otherwise PGDG is used. This only affects new clusters.                                                                                                  |
| Change  | You must now choose a failover manager explicitly when running <i>tpaexec configure</i> with the M1 architecture.                                                                                                                             |
| Bug fix | Fixed an issue with creation of PGD subscriber-only nodes whereby TPA incorrectly required 'subscriber-only' to be set on the replica instead of the upstream instance.                                                                       |
| Bug fix | TPA will now skip inapplicable tasks when deploying to containers even if you are using the 'bare' platform option (previously these were skipped only if 'docker' was selected).                                                             |
| Bug fix | Fixed an issue with permissions on <i>/etc/edb</i> whereby if you added the pgd-proxy role to a data node in a deployed PGD5 cluster, pgd-proxy would fail to start because it did not have permissions to open <i>pgd-proxy-config.yml</i> . |
| Bug fix | Fixed an issue whereby <i>/var/log/postgres</i> could end up with inappropriate permissions (0600) if a strict umask was set                                                                                                                  |
| Bug fix | Fixed an issue whereby repeating <i>tpaexec deploy</i> on a Barman instance correctly registered with PEM would lose the PEM Agent Barman configuration.                                                                                      |

### 1.2.25 Trusted Postgres Architect 23.20 release notes

Released: 01 Aug 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.20 include the following:

| Type        | Description                                                                                                                                                                                         |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| New feature | TPA now supports upgrades from PGD 4 to PGD 5 by running the new command <i>tpaexec reconfigure</i> to generate a revised <i>config.yml</i> and then <i>tpaexec upgrade</i> to perform the upgrade. |
| Enhancement | Added a new subcommand <i>tpaexec info validate</i> that runs a checksum over the TPA installation and confirms that it matches the one distributed with the package.                               |
| Change      | The <i>update-postgres</i> command has been replaced with the more general <i>upgrade</i> command.                                                                                                  |
| Change      | TPA now explicitly adds <i>tzdata-java</i> when installing OpenJDK for Failover Manager on RHEL 8 or 9. This is a workaround for <a href="#">this OpenJDK bug</a> .                                 |
| Change      | TPA now uses the latest available Debian AMIs on AWS (latest at the time of this release).                                                                                                          |
| Change      | TPA now runs <i>tpaexec provision</i> automatically as part of <i>tpaexec deploy</i> or <i>tpaexec upgrade</i> if <i>config.yml</i> has changed.                                                    |
| Bug fix     | Fixed a bug whereby TPA could attempt to use a non-existent user when running <i>pgd-cli</i> on <i>pgd-proxy</i> nodes.                                                                             |
| Bug fix     | Fixed a bug whereby changes made by <i>tpaexec relink</i> were not committed to the Git repository correctly.                                                                                       |

### 1.2.26 Trusted Postgres Architect 23.19 release notes

Released: 12 Jul 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.19 include the following:

| Type        | Description                                                                                                                                                                                                                                                                          |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| New feature | TPA now allows the creation of physical replicas of subscriber-only PGD nodes.                                                                                                                                                                                                       |
| New feature | TPA now supports the configuration of HTTP(S) HARP and PGD Proxy health probes.                                                                                                                                                                                                      |
| New feature | TPA now allows you to select Patroni as a failover manager with the M1 architecture. This support is experimental and not yet recommended for use in production.                                                                                                                     |
| Enhancement | TPA now allows you to set specific versions for edb-pgd-proxy and edb-bdr-utilities rather than always using the latest version.                                                                                                                                                     |
| Change      | On Debian-like systems, the package selection code now uses <code>-dbg</code> rather than <code>-dbgsym</code> for certain packages where applicable.                                                                                                                                |
| Change      | When configuring replication slots, TPA will now ensure that only valid characters are used in the <code>primary_slot_name</code> . Previously TPA would use the <code>inventory_hostname</code> as a default, which could contain hyphens; these are now replaced with underscores. |
| Change      | The default Failover Manager version is now 4.7.                                                                                                                                                                                                                                     |
| Bug fix     | Fixed an issue whereby PGD 3.7 to 4 upgrades would fail in TPA 23.18.                                                                                                                                                                                                                |
| Bug fix     | Fixed an issue whereby TPA would include underscores in TLS certificate Common Names when deploying PEM. This is invalid and would result in failure on some platforms.                                                                                                              |
| Bug fix     | Fixed an issue whereby an incorrect <code>etcd</code> service name would be used on Debian-like platforms, preventing TPA from starting <code>etcd</code> .                                                                                                                          |
| Bug fix     | Fixed an issue whereby TPA could not install <code>etcd</code> packages on RHEL 8.                                                                                                                                                                                                   |
| Bug fix     | Fixed an issue whereby the message <i>Failed to commit files to git: b''</i> would be displayed during configure.                                                                                                                                                                    |
| Bug fix     | Fixed an issue whereby TPA would erroneously generate and overwrite Postgres user passwords when <code>generate_password: false</code> .                                                                                                                                             |
| Bug fix     | Fixed an issue whereby volume map creation on AWS failed to take account of region resulting in failures when using regions other than <code>eu-west-1</code> .                                                                                                                      |

### 1.2.27 Trusted Postgres Architect 23.18 release notes

Released: 23 May 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.18 include the following:

| Type        | Description                                                                                                     |
|-------------|-----------------------------------------------------------------------------------------------------------------|
| Enhancement | TPA now uses <code>pg_basebackup</code> instead of <code>repmgr</code> for initial replica creation.            |
| Enhancement | TPA now supports SLES 15, excluding creation of local repositories for air-gapped deployments.                  |
| Enhancement | TPA now supports minor-version upgrades of PGD5.                                                                |
| Enhancement | TPA now runs improved tests when <code>tpaexec test</code> is executed.                                         |
| Bug fix     | Fixed an issue whereby TPA attempted to use legacy 2ndQuadrant repositories on unsupported distributions.       |
| Bug fix     | Fixed an issue whereby TPA didn't install <code>pg_receivewal</code> on Barman instances where it was required. |

continues on next page

Table 42 – continued from previous page

| Type    | Description                                                                                                                                    |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------|
| Bug fix | Fixed an issue whereby TPA intermittently failed to create symlinks to block devices on AWS hosts during provisioning, causing deploy to fail. |

## 1.2.28 Trusted Postgres Architect 23.17 release notes

Released: 10 May 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.17 include the following:

| Type        | Description                                                                                                                                                                                                                                                                                                                                                                   |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enhancement | Added a new <code>--pgd-proxy-routing</code> parameter to the <code>configure</code> command. This can be set to <i>global</i> or <i>local</i> . Local routing will make every PGD-Proxy route to a write leader within its own location. Global routing will make every proxy route to a single write leader, elected amongst all available data nodes across all locations. |
| Change      | Removed the <code>--active-locations</code> parameter from the <code>configure</code> command.                                                                                                                                                                                                                                                                                |
| Enhancement | TPA now supports Ubuntu 22.04                                                                                                                                                                                                                                                                                                                                                 |
| Change      | Updated the AWS AMIs used for RHEL 7 and 8.                                                                                                                                                                                                                                                                                                                                   |
| Bug fix     | Fixed an issue whereby TPA would incorrectly remove groups from existing Postgres users.                                                                                                                                                                                                                                                                                      |
| Bug fix     | Fixed an issue whereby TPA would print an unhelpful error message when a git commit failed.                                                                                                                                                                                                                                                                                   |
| Bug fix     | Fixed an issue whereby group names were incorrectly sanitized and uppercase letters were converted to underscores rather than lowercase ones.                                                                                                                                                                                                                                 |
| Bug fix     | Fixed an issue whereby Postgres was not restarted when required after CAMO configuration.                                                                                                                                                                                                                                                                                     |
| Bug fix     | Fixed an issue with <code>etcd</code> changes, ensuring that they are now idempotent and avoiding unnecessary restarts of <code>etcd</code> on subsequent deployments.                                                                                                                                                                                                        |

## 1.2.29 Trusted Postgres Architect 23.16 release notes

Released: 21 Mar 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.16 include the following:

| Type        | Description                                                                                                                                                                                                                                         |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Change      | The default PGD-Always-ON cluster is now one location with an associated subgroup containing two data nodes and one witness node.                                                                                                                   |
| Change      | TPA now deploys <code>pgd-proxy</code> on all data nodes by default.                                                                                                                                                                                |
| Enhancement | Added a new option, <code>--add-proxy-nodes-per-location N</code> , which creates separate proxy instances                                                                                                                                          |
| Enhancement | TPA now adds a witness node automatically if <code>--data_nodes_per_location</code> is even and prints a warning if you specify a cluster with only two locations                                                                                   |
| Change      | The parameter <code>--add-witness-only-location</code> has been renamed to <code>--witness-only-location</code> because we're NOT adding a location, but designating an already-named (in <code>--location-names</code> ) location as witness-only. |
| Change      | You must now specify Postgres flavour and version explicitly at <code>tpaexec configure</code> time                                                                                                                                                 |

continues on next page

Table 44 – continued from previous page

| Type        | Description                                                                                                           |
|-------------|-----------------------------------------------------------------------------------------------------------------------|
| Enhancement | Added new CLI abbreviations for Postgres flavour and version, for example <i>--postgresql 14</i> or <i>--edbpg 15</i> |
| Enhancement | Improved handling and documentation of the various supported EDB software repositories                                |
| Change      | TPA no longer includes the PGDG repository by default for PGD-Always-ON clusters                                      |
| Bug fix     | Fixed an issue whereby EDB Failover Manager was not selected as the failover manager for EPAS by default              |
| Bug fix     | Fixed an issue whereby pglogical was unnecessarily installed in the M1 architecture                                   |

### 1.2.30 Trusted Postgres Architect 23.15 release notes

Released: 15 Mar 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.15 include the following:

| Type         | Description                     |
|--------------|---------------------------------|
| Minor change | Changes to dependency mappings. |

### 1.2.31 Trusted Postgres Architect 23.14 release notes

Released: 23 Feb 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.14 include the following:

| Type    | Description                                                                                                |
|---------|------------------------------------------------------------------------------------------------------------|
| Bug fix | Fixes an error whereby package lists weren't correctly populated for PGD 3 and 4 configurations. (TPA-365) |
| Change  | Use multi-line BDR DCS configuration in HARP's config.yaml (TPA-360, RT90034)                              |

### 1.2.32 Trusted Postgres Architect 23.13 release notes

Released: 22 Feb 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.13 include the following:

| Type    | Description                                                      |
|---------|------------------------------------------------------------------|
| Bug fix | Don't enable old EDB repo with PGD-Always-ON and <i>--epas</i> . |
| Bug fix | Fix error with PGD-Always-ON and <i>--postgres-version 15</i> .  |

### 1.2.33 Trusted Postgres Architect 23.12 release notes

Released: 21 Feb 2023

New features, enhancements, bug fixes, and other changes in Trusted Postgres Architect 23.12 include the following:

| Type        | Description                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Feature     | Introduce full support for EDB Postgres Distributed 5, including Commit At Most Once (CAMO) configuration support based on commit scopes.                                                                                                                                                                                                                                                                          |
| Feature     | Introduce support for EDB Postgres Extended repository and packages.                                                                                                                                                                                                                                                                                                                                               |
| Enhancement | <p>&lt;p&gt;Preliminary support for configuring multi-region AWS clusters.&lt;/p&gt;&lt;p&gt;Multi-region clusters require manual setup of VPCs and VPC.&lt;/p&gt;</p>                                                                                                                                                                                                                                             |
| Enhancement | Enable proxy routing (and, therefore, subgroup RAFT) automatically for <i>--active-locations</i> . Removes the configure option to enable subgroup RAFT globally.                                                                                                                                                                                                                                                  |
| Bug fix     | Ensure the EDB_SUBSCRIPTION_TOKEN is not logged.                                                                                                                                                                                                                                                                                                                                                                   |
| Bug fix     | Allow the user to suppress addition of the products/default/release repo to tpa_2q_repositories.                                                                                                                                                                                                                                                                                                                   |
| Bug fix     | <p>&lt;p&gt;Ensure that nodes subscribe to bdr_child_group, if available.&lt;/p&gt;&lt;p&gt; In clusters with multiple subgroups, TPA did not expect instances to be subscribed to the replication sets for both the top group and the subgroup, so it would incorrectly remove the latter from the node's subscribed replication sets.&lt;/p&gt;</p>                                                              |
| Bug fix     | <p>&lt;p&gt;Fail reliably with a useful error if Postgres doesn't start. &lt;/p&gt;&lt;p&gt; Due to an Ansible bug, the deployment wouldn't fail if Postgres did not start on some instances, but did start on others (for example, due to a difference in the configuration). Continuing on with the deployment resulted in errors when trying to access cluster_facts for the failed hosts later. &lt;/p&gt;</p> |
| Bug fix     | <p>&lt;p&gt;Don't call <i>bdr.alter_node_replication_sets()</i> on witnesses for BDR 4.3 and later.&lt;/p&gt;&lt;p&gt; This adjusts to a new restriction in BDR versions where witness nodes are not handled with a custom replication set configuration. &lt;/p&gt;</p>                                                                                                                                           |
| Bug fix     | Replace hardcoded "barman" references to enable use of the barman_{user.group} settings to customize the barman user and home directory.                                                                                                                                                                                                                                                                           |
| Bug fix     | Add shared_preload_libraries entries, where appropriate, for extensions mentioned under postgres_databases[*].extensions.                                                                                                                                                                                                                                                                                          |
| Bug fix     | Ensure that <i>pgaudit</i> does not appear before <i>bdr</i> in shared_preload_libraries (to avoid a known crash).                                                                                                                                                                                                                                                                                                 |
| Bug fix     | Fix syntax error (DSN quoting) in pgd-cli config file.                                                                                                                                                                                                                                                                                                                                                             |
| Bug fix     | <p>&lt;p&gt; Sort endpoints in pgd-proxy config to avoid file rewrites. &lt;/p&gt; &lt;p&gt; This will likely require a pgd-proxy restart on the next deploy (but it will avoid unnecessary future rewrites/restarts on subsequent deploys). &lt;/p&gt;</p>                                                                                                                                                        |
| Bug fix     | Fix an error while installing rsync from a local-repo on RH systems.                                                                                                                                                                                                                                                                                                                                               |
| Bug fix     | Fix an error with Apache WSGI module configuration for PEM 9 on Debian systems.                                                                                                                                                                                                                                                                                                                                    |
| Bug fix     | Don't remove the bdr extension if it has been created on purpose, even if it is unused.                                                                                                                                                                                                                                                                                                                            |

## 1.2.34 Trusted Postgres Architect 23.1 to 23.11 release notes

### TPA 23.11

Released: 2023-01-31

#### Notable changes

- TPA-180 Introduce experimental support for PGD-Always-ON architecture (to be released later this year). PGD-Always-ON architecture will use the upcoming BDR version 5. Initial support has been added for internal purposes and will be improved in upcoming releases.

## Minor changes

- TPA-349 Bump dependency versions Bump cryptography version from 38.0.4 to 39.0.0 Bump jq version from 1.3.0 to 1.4.0
- TPA-345 Change TPAexec references to TPA in documentation. Update the documentation to use ‘TPA’ instead of ‘TPAexec’ when referring to the product.

## TPA 23.10

Released: 2023-01-04

## Minor changes

- TPA-161 Introduce `harp_manager_restart_on_failure` setting (defaults to false) to enable process restart on failure for the harp-manager systemd service

## Bug Fixes

- TPA-281 Delete FMS security groups when deprovisioning an AWS cluster Fixes a failure to deprovision a cluster’s VPC because of unremoved dependencies.
- TPA-305 Add `enterprisedb_password` to pre-generated secrets for Tower
- TPA-306 Prefer `PEM_PYTHON_EXECUTABLE`, if present, to `/usr/bin/python3` Fixes a Python module import error during deployment with PEM 9.0.
- TPA-219 Make pem-agent monitor the `bdr_database` by default on BDR instances

## TPA 23.9

Released: 2022-12-12

## Bugfixes

- TPA-301 Fix auto-detection of `cluster_dir` for Tower clusters When setting `cluster_dir` based on the Tower project directory, we now correctly check for the existence of the directory on the controller, and not on the instances being deployed to.
- TPA-283 Add dependency on `psutil`, required for Ansible Tower.
- TPA-278 Remove “`umask 0`” directive from `rsyslog` configuration, which previously resulted in the creation of world-readable files such as `rsyslogd.pid`.
- TPA-291 Respect the `postgres_package_version` setting when installing the Postgres server package to obtain `pg_receivewal` on Barman instances.

## TPA 23.8

Released: 2022-11-30

## Notable changes

- TPA-18 Support Ansible Tower 3.8 This release supports execution of `deploy.yml` (only) on a bare cluster (i.e., with existing servers) through Ansible Tower 3.8. Install TPAexec on the Tower server and run `tpaexec setup` to create a virtual environment which can be used in Tower Templates to run TPAexec playbooks. Use the `--use-ansible-tower` and `--tower-git-repository` configure options to generate a Tower-compatible cluster configuration. For details, see [Ansible Tower](#).

### Minor changes

- TPA-238 Initialise the cluster directory as a git repository If git is available on the system where you run TPAexec, `tpaexec configure` will now initialise a git repository within the cluster directory by default. If git is not available, it will continue as before. To avoid creating the repository (for example, if you want to store the cluster directory within an existing repository), use the `--no-git` option.

### TPA 23.7

Released: 2022-11-09

### Notable changes

- TPA-234 Support the community release of Ansible 2.9 TPAexec used to require the 2ndQuadrant/ansible fork of Ansible 2.9. In this release, you may instead choose to use the community release of Ansible with the `tpaexec setup --use-community-ansible`. For now, the default continues to be to use 2ndQuadrant/ansible. This will change in a future release; support for 2ndQuadrant/ansible will be dropped, and Ansible will become the new default.

### Minor changes

- TPA-209 Accept `--postgres-version 15` as a valid `tpaexec configure` option, subsequent to the release of Postgres 15
- TPA-226 Accept IP addresses in the `--hostnames-from` file Formerly, the file passed to `tpaexec configure` was expected to contain one hostname per line. Now it may also contain an optional IP address after each hostname. If present, this address will be set as the `ip_address` for the corresponding instance in `config.yml`. (If you specify your own `--hostnames-from` file, the hostnames will no longer be randomised by default.)
- TPA-231 Add a new `bdr-pre-group-join` hook This hook is executed before each node joins the BDR node group. It may be used to change the default replication set configuration that TPAexec provides.
- TPA-130 Use the `postgresql_user` module from `community.postgresql` The updated module from the `community.postgresql` collection is needed in order to correctly report the task status when using a SCRAM password (the default module always reports `changed`).
- TPA-250 Upgrade to the latest versions of various Python dependencies

### Bugfixes

- TPA-220 Ensure `LD_LIBRARY_PATH` in `.bashrc` does not start with `“:”`
- TPA-82 Avoid removing BDR-internal `_${group_name}_ext` replication sets
- TPA-247 Fix `“‘str object’ has no attribute ‘node_dsn’”` errors on AWS The code no longer assigns `hostvars[hostname]` to an intermediate variable and expects it to behave like a normal dict later (which works only sometimes). This fixes a regression in 23.6 reported for AWS clusters with PEM enabled, but also fixes other similar errors throughout the codebase.
- TPA-232 Eliminate a race condition in creating a symlink to generated secrets in the inventory that resulted in `“Error while linking: [Errno 17] File exists”` errors
- TPA-252 Restore code to make all BDR nodes publish to the witness-only replication set This code block was inadvertently removed in the v23.6 release as part of the refactoring work done for TPA-193.

## **TPA 23.6**

Released: 2022-09-28

### **Notable changes**

- TPA-21 Use boto3 (instead of the unmaintained boto2) AWS client library for AWS deployments. This enables SSO login and other useful features.
- TPA-202 Add harp-config hook. This deploy-time hook executes after HARP is installed and configured and before it is started on all nodes where HARP is installed.

### **Bugfixes**

- TPA-181 Set default python version to 2 on RHEL 7. Formerly, tpaexec could generate a config.yml with the unsupported combination of RHEL 7 and python 3.
- TPA-210 Fix aws deployments using existing security groups. Such a deployment used to fail at provision-time but will now work as expected.
- TPA-189 Remove group\_vars directory on deprovision. This fixes a problem that caused a subsequent provision to fail because of a dangling symlink.
- TPA-175 Correctly configure systemd to leave shared memory segments alone. This only affects source builds.
- TPA-160 Allow version setting for haproxy and PEM. This fixes a bug whereby latest versions of packages would be installed even if a specific version was specified.
- TPA-172 Install EFM on the correct set of hosts. EFM should be installed only on postgres servers that are members of the cluster, not servers which have postgres installed for other reasons, such as PEM servers.
- TPA-113 Serialize PEM agent registration. This avoids a race condition when several hosts try to run pemworker --register-agent at the same time.

## **TPA 23.5**

Released: 2022-08-23

### **Notable changes**

- TPA-81 Publish tpaexec and tpaexec-deps packages for Ubuntu 22.04 Jammy
- TPA-26 Support harp-proxy and harp-manager installation on a single node. It is now possible to have both harp-proxy and harp-manager service running on the same target node in a cluster.

## **TPA 23.4**

Released: 2022-08-03

### **Bugfixes**

- TPA-152 fix an issue with locale detection during first boot of Debian instances in AWS Hosts would fail to complete first boot which would manifest as SSH key negotiation issues and errors with disks not found during deployment. This issue was introduced in 23.3 and is related to TPA-38

### TPA 23.3

Released: 2022-08-03

#### Notable changes

- TPA-118 Exposed two new options in harp-manager configuration. The first sets HARP `harp_db_request_timeout` similar to `dcg request_timeout` but for database connections and the second `harp_ssl_password_command` specifies a command used to de-obfuscate `sslpassword` used to decrypt the `sslkey` in SSL enabled database connection

#### Minor changes

- TPA-117 Add documentation update on the use of wildcards in `package_version` options in `tpaexec config.yml`. This introduces a warning that unexpected package upgrades can occur during a `deploy` operation. See documentation in `tpaexec-configure.md` for more info
- TPA-38 Add locale files for all versions of Debian, and RHEL 8 and above. Some EDB software, such as Barman, has a requirement to set the user locale to `en_US.UTF-8`. Some users may wish to also change the locale, character set or language to a local region. This change ensures that OS files provided by `libc` are installed on AWS instances during firstboot using user-data scripts. The default locale is `en_US.UTF-8`. See `platform_aws.md` documentation for more info
- TPA-23 Add log config for syslog for cluster services Barman, HARP, repmgr, PgBouncer and EFM. The designated log server will store log files received in `/var/log/hosts` directories for these services
- TPA-109 Minor refactoring of the code in `pgbench` role around choosing lock timeout syntax based on a given version of BDR

#### Bugfixes

- TPA-147 For clusters that use the source install method some missing packages for Debian and Rocky Linux were observed. Debian receives library headers for `krb5` and `lz4`. On RedHat derived OSs the mandatory packages from the “Development Tools” package group and the `libcurl` headers have been added
- TPA-146 Small fix to the method of package selection for clusters installing Postgres 9.6
- TPA-138 Addresses a warning message on clusters that use the “bare” platform that enable the local-repo configure options. As the OS is not managed by TPAexec in the bare platform we need to inform the user to create the local-repo structure. This previously caused an unhandled error halting the configure progress
- TPA-135 When using `--use-local-repo-only` with the “docker” platform and the Rocky Linux image initial removal of existing yum repository configuration on nodes would fail due to the missing commands `find` and `xargs`. This change ensures that if the `findutils` package exists in the source repo it will be installed first
- TPA-111 Remove a redundant additional argument on the command used to register agents with the PEM server when `--enable-pem` option is given. Previously, this would have caused no problems as the first argument, the one now removed, would be overridden by the second
- TPA-108 Restore SELinux file context for postmaster symlink when Postgres is installed from source. Previously, a cluster using a SELinux enabled OS that is installing postgres from source would fail to restart Postgres as the `systemd` daemon would be unable to read the symlink stored in the Postgres data bin directory. This was discovered in tests using a recently adopted Rocky Linux image in AWS that has SELinux enabled and in enforcing mode by default

## TPA 23.2

Released: 2022-07-13

### Notable changes

- Add support for Postgres Backup API for use with Barman and PEM. Accessible through the `--enable-pg-backup-api` option.
- SSL certificates can now be created on a per-service basis, for example the server certificate for Postgres Backup API proxy service. Certificates will be placed in `/etc/tpa/<service>/<hostname>.cert`. These certificates can also be signed by a CA certificate generated for the cluster.
- Placement of Etcd for the BDR-Always-ON architecture When using `'harp_consensus_protocol: etcd'`, explicitly add `'etcd'` to the role for each of the following instances:
  - BDR Primary ('bdr' role)
  - BDR Logical Standby ('bdr' + 'readonly' roles)
  - only for the Bronze layout: BDR Witness ('bdr' + 'witness' roles)
  - only for the Gold layout: Barman ('barman' role) Credit: Gianni Ciolli [gianni.ciolli@enterprisedb.com](mailto:gianni.ciolli@enterprisedb.com)

### Minor changes

- Replace configure argument `--2q` with `--pgextended` to reflect product branding changes. Existing configuration will retain expected behaviour.
- Improve error reporting on Docker platform compatibility checks when using version 18 of docker, which comes with Debian old stable.
- Add some missing commands to CLI help documentation.
- Improved error reporting of configure command.
- Add initial support for building BDR 5 from source. Credit: Florin Irion [florin.irion@enterprisedb.com](mailto:florin.irion@enterprisedb.com)
- Changes to ensure ongoing compatibility for migration from older versions of Postgres with EDB products.

### Bugfixes

- Fixed an issue which meant packages for etcd were missing when using the `download-packages` command to populate the local-repo.
- Fixed an issue affecting the use of efm failover manager and the selection of its package dependencies

## TPA 23.1

Released: 2022-06-21

This release requires you to run `tpaexec setup` after upgrading (and will fail with an error otherwise)

### Changes to package installation behavior

In earlier versions, running `tpaexec deploy` could potentially upgrade installed packages, unless an exact version was explicitly specified (e.g., by setting `postgres_package_version`). However, this was never a safe, supported, or recommended way to upgrade. In particular, services may not have been safely and correctly restarted after a package upgrade during deploy.

With this release onwards, `tpaexec deploy` will never upgrade installed packages. The first deploy will install all required packages (either a specific version, if set, or the latest available), and subsequent runs will see that the package is installed, and do nothing further. This is a predictable and safe new default behavior.

If you need to update components, use `tpaexec update-postgres`. In this release, the command can update Postgres and Postgres-related packages such as BDR or pglogical, as well as certain other components, such as HARP, pgbouncer, and etcd (if applicable to a particular cluster). Future releases will safely support upgrades of more components.

### Notable changes

- Run “`harpcctl apply`” only if the HARP bootstrap config is changed **WARNING:** This will trigger a single harp service restart on existing clusters when you run `tpaexec deploy`, because `config.yml` is changed to ensure that lists are consistently ordered, to avoid unintended changes in future deploys
- Add `tpaexec download-packages` command to download all packages required by a cluster into a local-repo directory, so that they can be copied to cluster instances in airgapped/disconnected environments. See `air-gapped.md` and `local-repo.md` for details
- Require `--harp-consensus-protocol <etcd|bdr>` configure option for new BDR-Always-ON clusters TPAexec no longer supplies a default value here because the choice of consensus protocol can negatively affect failover performance, depending on network latency between data centres/locations, so the user is in a better position to select the protocol most suitable for a given cluster. This affects the configuration of newly-generated clusters, but does not affect existing clusters that use the former default of `etcd` without setting `harp_consensus_protocol` explicitly

### Minor changes

- Install `openjdk-11` instead of `openjdk-8` for EFM on distributions where the older version is not available
- Accept `harp_log_level` setting (e.g., under `cluster_vars`) to override the default harp-manager and harp-proxy log level (info)
- Configure harp-proxy to use a single multi-host BDR DCS endpoint DSN instead of a list of individual endpoint DSNs, to improve resilience
- Omit extra connection attributes (e.g., `ssl*`) from the local (Unix socket) DSN for the BDR DCS for harp-manager

### Bugfixes

- Ensure that harp-manager and harp-proxy are restarted if their config changes
- Fix harp-proxy errors by granting additional (new) permissions required by the readonly `harp_dcs_user`
- Disable BDR4 transaction streaming when CAMO is enabled If `bdr.enable_camo` is set, we must disable `bdr.default_streaming_mode`, which is not compatible with CAMO-protected transactions in BDR4. This will cause a server restart on CAMO-enabled BDR4 clusters (which could not work with streaming enabled anyway).

## 1.3 TPA installation

To use TPA, you need to install from packages or source and run the `tpaexec setup` command. This document explains how to install TPA packages. If you have an EDB subscription plan, and therefore have access to the EDB repositories, you should follow these instructions. To install TPA from source, please refer to [Installing TPA from source](#).

See [Distribution support](#) for information on what platforms are supported.

**Note**

Please make absolutely sure that your system has the correct date and time set, because various things will fail otherwise. We recommend you use a network time, for example `sudo ntpdate pool.ntp.org`

### 1.3.1 Quickstart

First, [subscribe to an EDB repository](#) .

#### Install TPA

```
sudo <your-package-manager> install tpaexec
```

#### Install additional dependencies

```
sudo /opt/EDB/TPA/bin/tpaexec setup
```

#### Verify installation (run as a normal user)

```
/opt/EDB/TPA/bin/tpaexec selftest
```

More detailed explanations of each step are given below.

### 1.3.2 Where to install TPA

As long as you are using a supported platform, TPA can be installed and run from your workstation. This is fine for learning, local testing or demonstration purposes. TPA supports [Installing Docker](#)

should you wish to perform a complete deployment on your own workstation.

For production use, we recommend running TPA on a dedicated, persistent virtual machine. We recommend this because it ensures that the cluster directories are retained and available to your team for future cluster management or update. It also means you only have to update one copy of TPA and you only need to provide network access from a single TPA host to the target instances.

### 1.3.3 Installing TPA packages

To install TPA, you must first [subscribe to an EDB repository](#) . TPA is available in all EDB repositories.

Install TPA as follows:

#### Install on Debian or Ubuntu

```
sudo apt-get install tpaexec
```

#### Install on RHEL, Rocky, AlmaLinux or Oracle Linux

```
sudo dnf install tpaexec
```

## SLES

```
sudo zypper install tpaexec
```

This will install TPA into `/opt/EDB/TPA`. It will also ensure that other required packages (e.g., Python 3.12 or later) are installed.

We mention `sudo` here only to indicate which commands need root privileges. You may use any other means to run the commands as root.

### 1.3.4 Setting up the TPA Python environment

Next, run `tpaexec setup` to create an isolated Python environment and install the correct versions of all required modules.

#### Note

On Ubuntu versions prior to 20.04, please use `sudo -H tpaexec setup` to avoid subsequent permission errors during *tpaexec configure*

```
sudo /opt/EDB/TPA/bin/tpaexec setup
```

You must run this as root because it writes to `/opt/EDB/TPA`, but the process will not affect any system-wide Python modules you may have installed (including Ansible).

Add `/opt/EDB/TPA/bin` to the `PATH` of the user who will normally run `tpaexec` commands. For example, you could add this to your `.bashrc` or equivalent shell configuration file:

```
export PATH=$PATH:/opt/EDB/TPA/bin
```

### 1.3.5 Installing TPA without internet or network access (air-gapped)

This section describes how to install TPA onto a server which cannot access either the EDB repositories, a Python package index, or both. For information on how to use TPA in such an environment, please see [Managing clusters in a disconnected or air-gapped environment](#)

#### Downloading TPA packages

If you cannot access the EDB repositories directly from the server on which you need to install TPA, you can download the packages from an internet-connected machine and transfer them. There are several ways to achieve this.

If your internet-connected machine uses the same operating system as the target, we recommend using `yumdownloader` (RHEL-like) or `apt download` (Debian-like) to download the packages.

Alternatively, you can download packages for any platform from your browser by visiting [EDB Repos](#) and selecting either ‘Enterprise’, ‘Standard’ or ‘Community 360’ under the heading ‘Download EDB software packages from your browser’. To install TPA you need these packages:

- `tpaexec`
- `tpaexec-deps`
- `edb-python312`

Once you have transferred the downloaded packages to the target server, you must install them using the appropriate tool for your platform.

## Installing without access to a Python package index

When you run `tpaexec setup`, it will ordinarily download the Python packages from a Python package index. Unless your environment provides a different index the default is the official [PyPI](#). If no package index is available, you should install the `tpaexec-deps` package in the same way you installed `tpaexec`. The `tpaexec-deps` package (available from the same repository as `tpaexec`) bundles everything that would have been downloaded, so that they can be installed without network access. Just install the package before you run `tpaexec setup` and the bundled copies will be used automatically.

### 1.3.6 Verifying your TPA installation

Once you're done with all of the above steps, run the following command to verify your local installation:

```
tpaexec selftest
```

If that command completes without any errors, your TPA installation is ready for use.

### 1.3.7 Upgrading TPA

To upgrade to a later release of TPA, you must:

1. Install the latest `tpaexec` package
2. Install the latest `tpaexec-deps` package (if required; see above)
3. Run `tpaexec setup` again

If you have subscribed to the TPA package repository as described above, running `apt-get update && apt-get upgrade` or `yum update` should install the latest available versions of these packages. If not, you can install the packages by any means available.

We recommend that you run `tpaexec setup` again whenever a new version of `tpaexec` is installed. Some new releases may not strictly require this, but others will not work without it.

### 1.3.8 Ansible versions

TPA uses `ansible-core 2.16` by default.

## 1.4 Open source TPA

### 1.4.1 What is Trusted Postgres Architect (TPA)?

TPA is an orchestration tool developed by [EnterpriseDB \(EDB\)](#) that uses Ansible to deploy Postgres clusters according to EDB's recommendations.

TPA embodies the best practices followed by EDB, informed by many years of hard-earned experience with deploying and supporting Postgres. These recommendations are as applicable to quick testbed setups as to production environments.

### 1.4.2 Next Steps

- [Installing TPA from source](#)
- [A First Cluster Deployment](#)
- [TPA's full documentation online](#)

### 1.4.3 TPA Open Source FAQs

#### Can I use this if I'm not an EDB customer?

Yes, TPA is an open source project under the GPLv3 license. It supports deploying clusters comprised of open source software, or EDB's proprietary products, or combinations.

#### Can I report an issue?

Yes, if you're an EDB customer then please contact support. Otherwise please open a GitHub Issue.

#### Can I contribute?

Sure, we'd love to hear from you but please open an issue before you start coding. We are quite selective with what TPA can/should do so bug fixes are more likely to get accepted than new features.

## 1.5 Installing TPA from source

This document explains how to use TPA from a copy of the source code repository.

#### Note

EDB customers must *TPA installation* in order to receive EDB support for the software.

To run TPA from source, you must install all of the dependencies (e.g., Python 3.12+) that the packages would handle for you, or download the source and *Running TPA in a Docker container*. (Either way will work fine on Linux and macOS.)

### 1.5.1 Quickstart

First, you must install the various dependencies Python 3, Python venv, git, openvpn and patch. Installing from EDB repositories would install these automatically along with the TPA packages.

Before you install TPA, you must install the required packages:

- **Debian/Ubuntu** `sudo apt-get install python3 python3-pip python3-venv git openvpn patch`
- **Redhat, Rocky or AlmaLinux (RHEL7)** `sudo yum install python3 python3-pip epel-release git openvpn patch`
- **Redhat, Rocky or AlmaLinux (RHEL8)** `sudo yum install python36 python3-pip epel-release git openvpn patch`

### 1.5.2 Clone and setup

With prerequisites installed, you can now clone the repository.

```
git clone https://github.com/enterprisedb/tpa.git ~/tpa
```

This creates a `tpa` directory in your home directory.

If you prefer to checkout with ssh use:

```
git clone ssh://git@github.com/EnterpriseDB/tpa.git ~/tpa
```

Add the bin directory, found within in your newly created clone, to your path with:

```
export PATH=$PATH:$HOME/tpa/bin
```

Add this line to your `.bashrc` file (or other profile file for your preferred shell).

You can now create a working tpa environment by running:

```
tpaexec setup
```

This will create the Python virtual environment that TPA will use in future. All needed packages are installed in this environment. To test this configured correctly, run the following:

```
tpaexec selftest
```

You now have tpaexec installed.

### 1.5.3 Dependencies

#### Python 3.12+

TPA requires Python 3.12 or later, available on most modern distributions. If you don't have it, you can use `pyenv` to install any version of Python you like without affecting the system packages.

```
# First, install pyenv and activate it in ~/.bashrc
# See https://github.com/pyenv/pyenv#installation
# (e.g., `brew install pyenv` on MacOS X)

pyenv install 3.12.0
```

```
Downloading Python-3.12.0.tar.xz...
- > https://www.python.org/ftp/python/3.12.0/Python-3.12.0.tar.xz
Installing Python-3.12.0...
Installed Python-3.12.0 to /home/ams/.pyenv/versions/3.12.0
```

```
pyenv local 3.12.0
pyenv version
```

```
3.12.0 (set by /home/ams/.pyenv/.python-version)
```

```
pyenv which python3
```

```
/home/ams/.pyenv/versions/3.12.0/bin/python3
```

```
python3 --version
```

```
3.12.0
```

If you were not already using `pyenv`, please remember to add `pyenv` to your `PATH` in `.bashrc` and call `eval "$ (pyenv init -) "` as described in the [pyenv documentation](#).

### Virtual environment options

By default, `tpaexec setup` will use the builtin `Python 3 -m venv` to create a `venv` under `$TPA_DIR/tpa-venv`, and activate it automatically whenever `tpaexec` is invoked.

You can run `tpaexec setup --venv /other/location` to specify a different location for the new `venv`.

We strongly suggest sticking to the default `venv` location. If you use a different location, you must also set the environment variable `TPA_VENV` to its location, for example by adding the following line to your `.bashrc` (or other shell startup scripts):

```
export TPA_VENV="/other/location"
```

### 2.1 A First Cluster Deployment

In this short tutorial, we are going to work through deploying a simple *MI* deployment onto a local Docker installation. By the end you will have four containers, one primary database, two replicas and a backup node, configured and ready for you to explore.

For this example, we will run TPA on an Ubuntu system, but the considerations are similar for most Linux systems.

#### 2.1.1 Installing TPA

If you're an EDB customer, you'll want to follow the *TPA installation* which will install the TPA packages straight from EDB's repositories.

If you are an open source user of TPA, there's *Installing TPA from source* which you can download from Github.com. Follow those guides and then return here.

#### 2.1.2 Installing Docker

As we said, We are going to deploy the example deployment onto Docker and unless you already have Docker installed we'll need to set that up.

On Debian or Ubuntu, install Docker by running:

```
sudo apt update
sudo apt install docker.io
```

For other Linux distributions, consult the [Docker Engine Install page](#) .

You will want to add your user to the docker group with:

```
sudo usermod -aG docker <yourusername>
newgrp docker
```

**Warning**

Giving a user the ability to speak to the Docker daemon lets them trivially gain root on the Docker host. Only trusted users should have access to the Docker daemon.

### 2.1.3 Creating a configuration with TPA

The next step in this process is to create a configuration. TPA does most of the work for you through its `configure` command. All you have to do is supply command line flags and options to select, in broad terms, what you want to deploy. Here's our `tpaexec configure` command:

```
tpaexec configure demo --architecture M1 --platform docker --postgresql 15 --enable-  
→repmgr --no-git
```

This creates a configuration called `demo` which has the *M1*. It will therefore have a primary, replica and backup node.

The `--platform docker` tells TPA that this configuration should be created on a local Docker instance; it will provision all the containers and OS requirements. Other platforms include [Using awscli to change volume attributes](#), which does the same with Amazon Web Services and [bare\(-metal servers\)](#), which skips to operating system provisioning and goes straight to installing software on already configured Linux hosts.

With `--postgresql 15`, we instruct TPA to use Community Postgres, version 15. There are several options here in terms of selecting software, but this is the most straightforward default for open-source users.

Adding `--enable-repmgr` tells TPA to use configure the deployment to use [Replication Manager](#) to hand replication and failover.

Finally, `--no-git` turns off the feature in TPA which allows you to revision control your configuration through git.

Run this command, and apparently, nothing will happen on the command line. But you will find a directory called `demo` has been created containing some files including a `config.yml` file which is a blueprint for our new deployment.

#### Provisioning the deployment

Now we are ready to create the containers (or virtual machines) on which we will run our new deployment. This can be achieved with the `provision` command. Run:

```
tpaexec provision demo
```

You will see TPA work through the various operations needed to prepare for deployment of your configuration.

#### Deploying

Once provisioned, you can move on to deployment. This installs, if needed, operating systems and system packages. It then installs the requested Postgres architecture and performs all the needed configuration.

```
tpaexec deploy demo
```

You will see TPA work through the various operations needed to deploy your configuration.

#### Testing

You can quickly test your newly deployed configuration using the `tpaexec test` command which will run `pgbench` on your new database.

```
tpaexec test demo
```

## Connecting

To get to a psql prompt, the simplest route is to log into one of the containers (or VMs or host depending on configuration) using docker or SSH. Run

```
tpaexec ping demo
```

to ping all the connectable hosts in the deployment: You will get output that looks something like:

```
unfair | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
uptake | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
quondam | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
uptight | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
```

Select one of the nodes which responded with `SUCCESS`. We shall use `uptake` for this example.

If you are only planning on using docker, use the command `docker exec -it uptake /bin/bash`, substituting in the appropriate hostname.

Another option, that works with all types of TPA deployment is to use SSH. To do that, first change current directory to the created configuration directory.

For example, our configuration is called `demo`, so we go to that directory. In there, we run `ssh -F ssh_config uptake` to connect.

```
cd demo
ssh -F ssh_config uptake
```

```
Last login: Wed Sep  6 10:08:01 2023 from 172.17.0.1
[root@uptake ~]#
```

In both cases, you will be logged in as a root user on the container.

We can now change user to the `postgres` user using `sudo -iu postgres`. As `postgres` we can run `psql`. TPA has already configured that user with a `.pgpass` file so there's no need to present a password.

```
sudo -iu postgres
psql
```

```
[root@uptake ~]# sudo -iu postgres
postgres@uptake:~ $ psql
psql (15.4)
Type "help" for help.
```

(continues on next page)

(continued from previous page)

```
postgres=#
```

And we are connected to our database.

You can connect from the host system without SSHing into one of the containers. Obtain the IP address of the host you want to connect to from the `ssh_config` file.

```
grep "^ *Host" demo/ssh_config
```

```
Host *
Host uptight
    HostName 172.17.0.9
Host unfair
    HostName 172.17.0.4
Host quondam
    HostName 172.17.0.10
Host uptake
    HostName 172.17.0.11
```

We are going to connect to uptake, so the IP address is 172.17.0.11.

You will also need to retrieve the password for the postgres user too. Run `tpaexec show-password demo postgres` to get the stored password from the system.

```
tpaexec show-password demo postgres
```

```
a9LmI1X^uMOpPoEnLuRdL%L$orQak3om
```

Assuming you have a Postgresql client installed, you can then run:

```
psql --host 172.17.0.11 -U postgres
```

```
Password for user postgres:
```

Enter the password you previously retrieved.

```
psql (14.9 (Ubuntu 14.9-0ubuntu0.22.04.1), server 15.4)
WARNING: psql major version 14, server major version 15.
        Some psql features might not work.
SSL connection (protocol: TLSv1.3, cipher: TLS_AES_256_GCM_SHA384, bits: 256,
↪compression: off)
Type "help" for help.

postgres=#
```

You are now connected from the Docker host to Postgres running in one of the TPA deployed Docker containers.

## 2.2 Cluster configuration

The `tpaexec configure` command generates a YAML cluster configuration file that is required by subsequent stages in the provision/deploy/test cycle.

## 2.2.1 Quickstart

```
tpaexec configure ~/clusters/speedy --architecture M1 \
  --postgresql 14 \
  --failover-manager repmgr
```

This command will create a directory named `~/clusters/speedy` and generate a configuration file named `config.yml` that follows the layout of the architecture named `M1` (single primary, `N` replicas). It will create a git repository in the new directory and make an initial commit containing the generated `config.yml`.

The command also accepts various options (some specific to the selected architecture or platform) to modify the configuration, but the defaults are sensible and intended to be usable straightaway. You are encouraged to read the generated `config.yml` and fine-tune the configuration to suit your needs. (Here's an overview of [Instance configuration](#).)

It's possible to write `config.yml` entirely by hand, but it's much easier to edit the generated file.

## 2.2.2 Configuration options

The first argument must be the cluster directory, e.g., `speedy` or `~/clusters/speedy` (the cluster will be named `speedy` in both cases). We recommend that you keep all your clusters in a common directory, e.g., `~/clusters` in the example above.

The next argument must be `--architecture <name>` to select an architecture, e.g., `M1` or `BDR-Always-ON` or `PGD-X`. For a complete list of architectures, run `tpaexec info architectures`.

Next, you must specify a *Postgres flavour and version* to install.

The arguments above are always mandatory. The rest of the options described here may be safely omitted, as in the example above; the defaults will lead to a usable result.

Run `tpaexec help configure-options` for a list of common options.

### Architecture-specific options

The architecture you select determines what other options are accepted. Typically, each architecture accepts some unique options as well as the generic options described below.

For example, with `M1` you can use `--location-names l1 l2` to create a cluster with nodes in two named locations. Please consult the documentation for an architecture for a list of options that it accepts (or, in some cases, requires).

### Platform options

Next, you may use `--platform <name>` to select a platform, e.g., *Using awscli to change volume attributes* or *bare(-metal servers)*.

An architecture may or may not support a particular platform. If not, it will fail to configure the cluster.

The choice of platform affects the interpretation of certain options. For example, if you choose `aws`, the arguments to `--region <region>` and `--instance-type <type>` must be a valid [AWS region name](#) and [EC2 instance type](#)

respectively. Please refer to the platform documentation for more details.

If you do not explicitly select a platform, the default is currently `aws`.

**Note:** TPA fully supports creating clusters with instances on different platforms, but `tpaexec configure` cannot currently generate such a configuration. You must edit `config.yml` to specify multiple platforms.

### Owner

Specify `--owner <name>` to associate the cluster (by some platform-specific means, e.g., AWS tags) with the name of a person responsible for it. This is especially important for cloud platforms. By default, the owner is set to the login name of the user running `tpaexec provision`.

(You may use your initials, or “Firstname Lastname”, or anything else that identifies you uniquely.)

### Region

Specify `--region <region>` to select a region.

This option is meaningful only for cloud platforms. The default for AWS is `eu-west-1`.

**Note:** TPA fully supports creating clusters that span multiple regions, but `tpaexec configure` cannot currently generate such a configuration. You must edit `config.yml` to specify multiple regions.

### Network configuration

!!!Note

These options are not meaningful for the “bare” platform, where TPA will not alter the network configuration of existing servers.

Specify `--network 192.168.0.0/16` to assign subnets from a different network. On AWS clusters, this corresponds to the VPC CIDR. See [Using awscli to change volume attributes](#) documentation for details.

Specify `--subnet-prefix 26` to assign subnets of a different size, /26 instead of /28 in this case.

!!!Note

When the “docker” platform is selected, TPA will always place the entire cluster in a single subnet regardless of the architecture. This subnet is generated according to the logic described here with the exception that if the `subnet-prefix` is not specified, TPA will automatically select a subnet size large enough to accommodate the number of instances in `config.yml`.

Specify `--exclude-subnets-from <directory>` to exclude subnets that are already used in existing cluster `config.yml` files. You can specify this argument multiple times for each directory.

### Instance type

Specify `--instance-type <type>` to select an instance type.

This option is meaningful only for cloud platforms. The default for AWS is `t3.micro`.

### Disk space

Specify `--root-volume-size 64` to set the size of the root volume in GB. (Depending on the platform, there may be a minimum size required for the root volume.)

The `--postgres-volume-size <size>` and `--barman-volume-size <size>` options are available to set the sizes of the Postgres and Barman volumes on those architectures and platforms that support separate volumes for Postgres and Barman.

None of these options is meaningful for the “bare” platform, where TPA has no control over volume sizes.

### Hostnames

By default, `tpaexec configure` will randomly select as many hostnames as it needs from a pre-approved list of several dozen names. This should be enough for most clusters.

Specify `--hostnames-from <filename>` to select hostnames from a file with one name per line. The file must contain at least as many valid hostnames as there are instances in your cluster. Each line may contain an optional IP address after the name; if present, this address will be set as the `ip_address` for the corresponding instance in `config.yml`. If two IP addresses are present, the first will be set as `public_ip` and the second as `private_ip`.

Use `--hostnames-pattern '...pattern...'` to limit the selection to lines matching an `egrep` pattern.

Use `--hostnames-sorted-by="--dictionary-order"` to select a `sort(1)` option other than `--random-sort` (which is the default).

Use `--hostnames-unsorted` to not sort hostnames at all. In this case, they will be assigned in the order they are found in the hostnames file. This is the default when a hostnames file is explicitly specified.

Use `--cluster-prefixed-hostnames` to make each hostname begin with the name of the cluster. This can be useful to avoid hostname clashes when running more than one docker cluster on the same host.

Hostnames may contain only letters (a-z), digits (0-9), and '-'. They may be FQDNs, depending on the selected platform. Hostnames should be in lowercase; any uppercase characters will be converted to lowercase internally, and any references to these hostnames in `config.yml` (e.g., `upstream: hostname`) must use the lowercase version.

## 2.2.3 Software selection

### Distribution

Specify `--distribution <name>` to select a distribution.

The selected platform determines which distributions are available, and which one is used by default.

In general, you should be able to use “Debian”, “RedHat”, “Ubuntu”, and “SLES” to select the right images.

Optionally, you can include the version of the distribution by passing `--os-version <version number>`.

This option is not meaningful for the “bare” platform, where TPA has no control over which distribution is installed.

### EDB repositories

TPA can enable any EDB software repository that you have access to through a subscription. By default, TPA will install any product repositories that the architecture requires.

More detailed explanation of how TPA uses EDB repositories is available [Configuring EDB Repos 2.0 repositories](#) and on the page for each architecture.

Specify `--edb-repositories repository ...` to specify the complete list of EDB repositories to install on each instance.

Use this option with care. TPA will configure the named repositories with no attempt to make sure the combination is appropriate.

To use this options, you must `export EDB_SUBSCRIPTION_TOKEN=xxx` before you run TPA. You can get an EDB token from [enterprisedb.com/repos](https://enterprisedb.com/repos).

### Local repository support

Use `--enable-local-repo` to create a local package repository from which to ship packages to target instances.

In environments with restricted network access, you can instead use `--use-local-repo-only` to create a local repository and disable all other package repositories on target instances, so that packages are installed only from the local repository.

The page about [Creating and using a local repository](#) has more details.

## Software versions

### Postgres flavour and version

TPA supports PostgreSQL, EDB Postgres Extended, and EDB Postgres Advanced Server (EPAS) versions 13 through 18.

You must specify both the flavour (or distribution) and major version of Postgres to install, for example:

- `--postgresql 14` will install PostgreSQL 14
- `--edb-postgres-extended 15` will install EDB Postgres Extended 15
- `--edb-postgres-advanced 15 --redwood` will install EPAS 15 in “Redwood” mode
- `--edb-postgres-advanced 15 --no-redwood` will install EPAS 15 in non-Redwood mode

If you are installing EPAS, you must specify whether it should operate in `--redwood` or `--no-redwood` mode, i.e., whether to enable or disable its Oracle compatibility features.

Installing EDB Postgres Extended or Postgres Advanced Server requires a valid [Configuring EDB Repos 2.0 repositories](#).

### Package versions

By default, we always install the latest version of every package. This is usually the desired behaviour, but in some testing scenarios, it may be necessary to select specific package versions using any of the following options:

1. `--postgres-package-version 10.4-2.pg90+1`
2. `--repmgr-package-version 4.0.5-1.pg90+1`
3. `--barman-package-version 2.4-1.pg90+1`
4. `--pglogical-package-version '2.2.0*'`
5. `--bdr-package-version '3.0.2*'`
6. `--pgbouncer-package-version '1.8*'`
7. `--beacon-agent-package-version 1.56.2-1`
8. `--etcd-package-version 9.8.0-1.el8`
9. `--patroni-package-version 4.0.0-1PGDG.rhel8` 10. `--pem-server-package-version 9.7.0-1.el9` 11. `--pem-agent-package-version 9.7.0-1.el9` 12. `--pg-backup-api-package-version 2.0.0-1.el8` 13. `--pgd-proxy-package-version 5.0.0-1` 14. `--pgdcli-package-version 5.6.1`

You may use any version specifier that apt or yum would accept.

If your version does not match, try appending a `*` wildcard. This is often necessary when the package version has an epoch qualifier like `2:...`.

You may also specify `--extra-packages p1 p2 ...` or `--extra-postgres-packages p1 p2 ...` to install additional packages. The former lists packages to install along with system packages, while the latter lists packages to install later along with postgres packages. (If you mention packages that depend on Postgres in the former list, the installation will fail because Postgres will not yet be installed.) The arguments are passed on to the package manager for installation without any modifications.

The `--extra-optional-packages p1 p2 ...` option behaves like `--extra-packages`, but it is not an error if the named packages cannot be installed.

### Known issue with wildcard use

Please note that the use of wildcards in `*_package_version` when added permanently to `config.yml`, can result in unexpected updates to installed software during `tpaexec deploy` on nodes with RHEL 8 and above (or derivative OSs which use `dnf` such as Rocky Linux). When `deploy` runs on an existing cluster that already has packages installed, `ansible` may be unable to match the full package version. For example, if the value for `bdr_package_version` was set to `3.6*` then `ansible` would not be able to match this to an installed version of PGD, it would assume no package is installed, and it would attempt to install the latest version available of the package with the same name in the configured repository, e.g. 3.7.

We are aware of this limitation as an `ansible dnf` module bug and hope to address this in a future release of TPA.

### Building and installing from source

If you specify `--install-from-source postgres`, Postgres will be built and installed from a git repository instead of installed from packages. By default, this will build the appropriate `REL_nnn_STABLE` branch.

You may use `--install-from-source postgres bdr5` to build and install both components from source, or just use `--install-from-source bdr5` to use packages for Postgres, but build and install PGD v5 from source. By default, this will build the `main` branch of PGD.

To build a different branch, append `:branchname` to the corresponding argument. For example `--install-from-source 2ndqpostgres:dev/xxx`, or `pglogical:bug/nnnn`.

You may not be able to install packages that depend on a package that you chose to replace with a source installation instead. For example, PGD v3 packages depend on `pglogical v3` packages, so you can't install `pglogical` from its source repository and PGD from packages. Likewise, you can't install Postgres from source and `pglogical` from packages.

## 2.2.4 Overrides

You may optionally specify `--overrides-from a.yml ...` to load one or more YAML files with settings to merge into the generated `config.yml`.

Any file specified here is first expanded as a Jinja2 template, and the result is loaded as a YAML data structure, and merged recursively into the arguments used to generate `config.yml` (comprising architecture and platform defaults and arguments from the command-line). This process is repeated for each additional override file specified; this means that settings defined by one file will be visible to any subsequent files.

For example, your override file might contain:

```
cluster_tags:
  some_tag: "{{ lookup(env, SOME_ENV_VAR) }}"

cluster_vars:
  synchronous_commit: remote_write
  postgres_conf_settings:
    effective_cache_size: 4GB
```

These settings will augment `cluster_tags` and `cluster_vars` that would otherwise be in `config.yml`. Settings are merged recursively, so `cluster_tags` will end up containing both the default `Owner` tag as well as `some_tag`. Similarly, the `effective_cache_size` setting will override that variable, leaving other `postgres_conf_settings` (if any) unaffected. In other words, you can set or override specific subkeys in `config.yml`, but you can't empty or replace `cluster_tags` or any other hash altogether.

The merging only applies to hash structures, so you cannot use this mechanism to change the list of `instances` within `config.yml`. It is most useful to augment `cluster_vars` and `instance_defaults` with common settings for your environment.

That said, the mechanism does not enforce any restrictions, so please exercise due caution. It is a good idea to generate two configurations with and without the overrides and diff the two `config.yml` files to make sure you understand the effect of all the overrides.

## 2.2.5 Ansible Tower

Use the `--use-ansible-tower` and `--tower-git-repository` options to create a cluster adapted for deployment with Ansible Tower. See *TPA and Ansible Tower/Ansible Automation Platform* for details.

## 2.2.6 PGAI agent

Use the `--enable-beacon-agent` and `--beacon-agent-project-id` options to install the PGAI agent (packaged as `beacon-agent`), which enables you to view your cluster in the EDB Postgres AI Console. See *Configuring the PGAI agent* for details.

## 2.2.7 Git repository

By default, a git repository is created with an initial branch named after the cluster, and a single commit is made, with the configure options you used in the commit message. If you don't have git in your `$PATH`, `tpaexec` will not raise an error but the repository will not be created. To suppress creation of the git repository, use the `--no-git` option. (Note that in an Ansible Tower cluster, a git repository is required and will be created later by `tpaexec provision` if it does not already exist.)

## 2.2.8 Keyring backend for vault password

TPA generates a cluster specific ansible vault password. This password is used to encrypt other sensitive variables generated for the cluster, postgres user password, barman user password and so on.

Keyring backend `system` will leverage the best keyring backend on your system from the list of supported backend by python keyring module including `gnome-keyring` and `secret-tool`.

Default is to store the vault password using `system` keyring for new cluster. removing `keyring_backend: system` in `config.yml` file **before** any `provision` will revert previous default to store vault password in plaintext file.

Using `keyring_backend: system` also generates a `vault_name` entry in `config.yml` used to store the vault password unique storage name. TPA generate an UUID by default but there is no naming scheme requirements.

Note: When using `keyring_backend: system` and the same base `config.yml` file for multiple clusters with same `cluster_name`, by copying the config file to a different location, ensure the value pair (`vault_name`, `cluster_name`) is unique for each cluster copy.

Note: When using `keyring_backend: system` and moving an already provisioned cluster folder to a different tpa host, ensure that you export the associated vault password on the new machine's system keyring. vault password can be displayed via `tpaexec show-vault <cluster_dir>`.

## 2.2.9 Security standards compliance

Use the `--compliance stig` or `--compliance cis` options to generate a cluster with configuration suitable for complying with the STIG or CIS standard. See *Security standards compliance* for details. Note that these options do not guarantee that the cluster fulfills the relevant standard; they only cause TPA to generate a configuration designed to comply with those aspects of the standard that can be controlled by TPA.

## 2.2.10 Examples

Let's see what happens when we run the following command:

```
tpaexec configure ~/clusters/speedy --architecture M1 \
  --distribution Debian \
  --os-version 12 \
  --platform aws --region us-east-1 --network 10.33.0.0/16 \
  --instance-type t2.medium --root-volume-size 32 \
  --postgres-volume-size 64 --barman-volume-size 128 \
  --postgresql 14 \
  --failover-manager repmgr
```

There is no output, so there were no errors. The cluster directory has been created and populated.

```
ls -lh ~/clusters/speedy/
```

```
total 8.0K
drwxrwxr-x 2 haroon haroon 4.0K Aug 17 02:33 commands
-rw-rw-r-- 1 haroon haroon 1.5K Aug 17 02:33 config.yml
lrwxrwxrwx 1 haroon haroon 53 Aug 17 02:33 deploy.yml -> /home/haroon/tpa/
↳ architectures/M1/deploy.yml
```

The cluster configuration is in config.yml, and its neighbours are links to architecture-specific support files that you need not interact with directly. Here's what the configuration looks like:

```
--
architecture: M1
cluster_name: speedy
cluster_tags: {}

keyring_backend: system
vault_name: cfae3da3-ec00-46cd-ab05-e153f1c788db

cluster_rules:
- cidr_ip: 0.0.0.0/0
  from_port: 22
  proto: tcp
  to_port: 22
- cidr_ip: 10.33.120.80/28
  from_port: 0
  proto: tcp
  to_port: 65535
ec2_ami:
  Name: debian-11-amd64-20240104-1616
  Owner: 136693071363
ec2_instance_reachability: public
ec2_vpc:
  us-east-1:
    Name: Test
    cidr: 10.33.0.0/16
```

(continues on next page)

(continued from previous page)

```
cluster_vars:
  edb_repositories: []
  failover_manager: repmgr
  postgres_flavour: postgresql
  postgres_version: 14
  preferred_python_version: python3

locations:
- Name: main
  az: us-east-1a
  region: us-east-1
  subnet: 10.33.120.80/28

instance_defaults:
  default_volumes:
  - device_name: root
    encrypted: true
    volume_size: 32
    volume_type: gp2
  - device_name: /dev/sdf
    encrypted: true
    vars:
      volume_for: postgres_data
      volume_size: 64
      volume_type: gp2
  platform: aws
  type: t2.medium
  vars:
    ansible_user: admin

instances:
- Name: uproar
  backup: kinsman
  location: main
  node: 1
  role:
  - primary
- Name: unravel
  location: main
  node: 2
  role:
  - replica
  upstream: uproar
- Name: kinsman
  location: main
  node: 3
  role:
  - barman
```

(continues on next page)

(continued from previous page)

```

- log-server
- witness
upstream: uproar
volumes:
- device_name: /dev/sdf
  encrypted: true
  vars:
    volume_for: barman_data
  volume_size: 128
  volume_type: gp2

```

The next step is to run *tpaexec provision*

or learn more about how to customise the configuration of *Cluster configuration* or how to configure an *Instance configuration*.

## 2.3 tpaexec deploy

Deployment is the process of installing and configuring Postgres and other software on the cluster's servers. This includes setting up replication, backups, and so on.

At the end of the deployment stage, Postgres will be up and running along with other components like repmgr, Barman, pgbouncer, etc. (depending on the architecture selected).

### 2.3.1 Prerequisites

Before you can run *tpaexec deploy*, you must have already run *Cluster configuration* to generate the cluster configuration and then provisioned the servers with *tpaexec provision*.

Before deployment, you must export `EDB_SUBSCRIPTION_TOKEN=xxx` if you are using any *Configuring EDB Repos 2.0 repositories*. If you forget to do this, an error message will soon remind you.

### 2.3.2 Quickstart

```
tpaexec deploy ~/clusters/speedy -v
```

```

Using /opt/EDB/TPA/ansible/ansible.cfg as config file

PLAY [Basic initialisation and fact discovery] *****
...

PLAY [Set up TPA cluster nodes] *****
...

PLAY RECAP *****
zealot           : ok=281  changed=116  unreachable=0    failed=0
keeper           : ok=284  changed=96   unreachable=0    failed=0
quaver           : ok=260  changed=89   unreachable=0    failed=0
quavery          : ok=260  changed=88   unreachable=0    failed=0
quirk            : ok=262  changed=100  unreachable=0    failed=0

real            7m1.907s

```

(continues on next page)

(continued from previous page)

```
user      3m2.492s
sys       1m5.318s
```

This command produces a great deal of output and may take a long time (depending primarily on the latency between the host running `tpaexec` and the hosts in the cluster, as well as how long it takes the instances to download the packages they need to install). We recommend that you use at least one `-v` during deployment. The output is also logged to `ansible.log` in the cluster directory.

The exact number of hosts, tasks, and changed tasks may of course vary.

The `deploy` command takes no options itself—any options you provide after the cluster name are passed on unmodified to Ansible (e.g., `-v`).

Those who are familiar with Ansible may be concerned by the occasional red “failed” task output scrolling by. Rest assured that if the process does not stop soon afterwards, the error is of no consequence, and the code will recover from it automatically.

When the deployment is complete, you can run `tpaexec test` to verify the installation.

### 2.3.3 Selective deployment

You can limit the deployment to a subset of your hosts by setting `deploy_hosts` to a comma-separated list of instance names:

```
tpaexec deploy ~/clusters/speedy -v -e deploy_hosts=keeper,quaver
```

This will run the deployment on the given instances, though it will also initially execute some tasks on other hosts to collect information about the state of the cluster.

(Setting `deploy_hosts` is the recommended alternative to using Ansible’s `--limit` option, which TPA does not support.)

### 2.3.4 deploy.yml

The deployment process is architecture-specific. Here’s an overview of the various *Instance configuration*. If you are familiar with Ansible playbooks, you can follow along as `tpaexec` applies various roles to the cluster’s instances.

Unlike `config.yml`, `deploy.yml` is not designed to be edited (and is usually a link into the architectures directory). Even if you want to extend the deployment process to run your own Ansible tasks, you should do so by *TPA hooks*. This protects you from future implementation changes within a particular architecture.

## 2.4 tpaexec provision

Provision creates instances and other resources required by the cluster.

The exact details of this process depend both on the architecture (e.g. *MI*) and platform (e.g. *Using awscli to change volume attributes*) that you selected while configuring the cluster.

At the end of the provisioning stage, you will have the required number of instances with the basic operating system installed, which TPA can access via `ssh` (with `sudo` to root).

### 2.4.1 Prerequisites

Before you can provision a cluster, you must generate the cluster configuration with *Cluster configuration* (and edit `config.yml` to fine-tune the configuration if needed).

You may need additional platform-dependent steps. For example, you need to obtain an AWS API access keypair to provision EC2 instances, or set up LXD or Docker to provision containers. Consult the platform documentation for details.

## 2.4.2 Quickstart

```
tpaexec provision ~/clusters/speedy
```

```
PLAY [Provision cluster] *****
...

TASK [Set up EC2 instances] *****
changed: [localhost] => (item=us-east-1:uproar)
changed: [localhost] => (item=us-east-1:unravel)
changed: [localhost] => (item=us-east-1:kinsman)
...

TASK [Generate ssh_config file for the cluster] *****
changed: [localhost]

PLAY RECAP *****
localhost                : ok=163  changed=35   unreachable=0    failed=0    _
↪skipped=44   rescued=0    ignored=2

real    4m42.726s
user    0m39.101s
sys     0m15.687s
```

This command will produce lots of output (append `-v`, `-vv`, etc. to the command if you want even more verbose output). The output is also logged to `ansible.log` in the cluster directory. This can be overridden by setting the environment variable `ANSIBLE_LOG_PATH` to the path and name of the desired logfile.

If it completes without error, you may proceed to run *tpaexec deploy* to install and configure software.

## 2.4.3 Options

When provisioning cloud instances, it is especially important to make sure instances are directly traceable to a human responsible for them. By default, TPA will tag EC2 instances as being owned by the login name of the user running `tpaexec provision`.

Specify `--owner <name>` to change the name (e.g., if your username happens to be something generic, like `postgres` or `ec2-user`). You may use initials, or “Firstname Lastname”, or anything else to uniquely identify a person.

Any other options you specify are passed on to Ansible.

## 2.4.4 Accessing the instances

After provisioning completes, you should be able to SSH to the instances (after a brief delay to allow the instances to boot up and install their SSH host keys). As shown in the output above, `tpaexec` will generate an `ssh_config` file for you to use.

```
cd ~/clusters/speedy
cat ssh_config
```

```
Host *
    Port 22
    IdentitiesOnly yes
    IdentityFile "id_speedy"
    UserKnownHostsFile known_hosts tpa_known_hosts
    ServerAliveInterval 60

Host uproar
    User admin
    HostName 3.88.255.205
Host unravel
    User admin
    HostName 54.80.99.142
Host kinsman
    User admin
    HostName 54.165.229.179
```

To login to a host, use the command `ssh -F ssh_config` followed by the hostname. For example `ssh -F ssh_config uproar`.

You can run `tpaexec deploy` immediately after provisioning. It will wait as long as required for the instances to come up. You do not need to wait for the instances to come up, or ssh in to them before you start deployment.

## 2.4.5 Generated files

During the provisioning process, a number of new files will be created in the cluster directory:

```
ls ~/clusters/speedy
```

```
total 240
-rw-r--r-- 1 ams ams 193098 Aug  4 17:59 ansible.log
drwxr-xr-x 2 ams ams  4096 Aug  4 17:38 commands
-rw-r--r-- 1 ams ams  1442 Aug  4 17:54 config.yml
lrwxrwxrwx 1 ams ams    51 Aug  4 17:38 deploy.yml ->
/opt/EDB/TPA/architectures/M1/deploy.yml
drwxr-xr-x 2 ams ams  4096 Aug  4 17:38 hostkeys
-rw----- 1 ams ams  1675 Aug  4 17:38 id_speedy
-rw----- 1 ams ams  1438 Aug  4 17:38 id_speedy.ppk
-rw-r--r-- 1 ams ams   393 Aug  4 17:38 id_speedy.pub
drwxr-xr-x 4 ams ams  4096 Aug  4 17:50 inventory
-rw-r--r-- 1 ams ams  2928 Aug  4 17:50 tpa_known_hosts
-rw-r--r-- 1 ams ams   410 Aug  4 17:50 ssh_config
-rw-r--r-- 1 ams ams  3395 Aug  4 17:59 vars.json
drwxr-xr-x 2 ams ams  4096 Aug  4 17:38 vault
```

We've already studied the *sshconfig file*, which refers to the `'id*' files` (an SSH keypair generated for the cluster) and `tpa_known_hosts` (the signatures of the `hostkeys/` installed on the instances).

The `vars.json` file may be used by `tpaexec provision` on subsequent invocations with `--cached`.

The `inventory/` directory contains static and dynamic inventory files as well as group and host variable definitions from `config.yml`.

```
cat inventory/00-speedy
```

```
[tag_Cluster_speedy]
uproar ansible_host=3.88.255.205 node=1 platform=aws
unravel ansible_host=54.80.99.142 node=2 platform=aws
kinsman ansible_host=54.165.229.179 node=3 platform=aws
```

```
cat inventory/group_vars/tag_Cluster_speedy/01-speedy.yml
```

```
cluster_name: speedy
cluster_tag: tag_Cluster_speedy
edb_repositories: []
failover_manager: repmgr
keyring_backend: system
postgres_flavour: postgresql
postgres_version: 14
preferred_python_version: python3
ssh_key_file: id_speedy
tpa_version: v23.33-24-g4c0909d1
```

```
cat inventory/host_vars/kinsman/01-instance_vars.yml
```

```
ansible_user: admin
location: main
region: us-east-1
role:

- barman

- log-server

- witness
upstream: uproar
volumes:

- device: /dev/xvda

- device: /dev/sdf
  volume_for: barman_data
```

If you now change a variable in config.yml and rerun provision, these files will be updated. If you don't change the configuration, it won't do anything. If you add a new instance in config.yml and rerun, it will bring up the new instance without affecting the existing ones.

## 2.5 tpaexec test

Now we run architecture-specific tests against a deployed cluster to verify the installation. At the end of this stage, we have a fully-functioning cluster.

You must have already run `tpaexec configure`, `tpaexec provision`, and `tpaexec deploy` successfully before you can run `tpaexec test`.

## 2.5.1 Quickstart

```
tpaexec test ~/clusters/speedy -v
```

Output is once again logged to `ansible.log` in the cluster directory.

If this command succeeds, your cluster works.

Congratulations.

## 3.1 PGD-Always-ON

!!!Note

This architecture is for Postgres Distributed 5 only. If you require PGD 4 or 3.7 please use *BDR-Always-ON* .

EDB Postgres Distributed 5 in an Always-ON configuration, suitable for use in test and production.

This architecture requires an EDB subscription. All software will be sourced from *Configuring EDB Repos 2.0 repositories* .

### 3.1.1 Cluster configuration

#### Overview of configuration options

An example invocation of `tpaexec configure` for this architecture is shown below.

```
tpaexec configure ~/clusters/pgd-ao \  
  --architecture PGD-Always-ON \  
  --edb-postgres-extended 15 \  
  --platform aws --instance-type t3.micro \  
  --distribution Debian \  
  --pgd-proxy-routing global \  
  --location-names dc1 dc2 dc3 \  
  --witness-only-location dc3 \  
  --data-nodes-per-location 2
```

You can list all available options using the help command.

```
tpaexec configure --architecture PGD-Always-ON --help
```

The table below describes the mandatory options for PGD-Always-ON and additional important options. More detail on the options is provided in the following section.

## Mandatory Options

| Options                                                           | Description                                    |
|-------------------------------------------------------------------|------------------------------------------------|
| <code>--architecture (-a)</code>                                  | Must be set to <i>PGD-Always-ON</i>            |
| Postgres flavour and version (e.g. <code>--postgresql 15</code> ) | A valid <i>Postgres flavour and version</i> .  |
| <code>--pgd-proxy-routing</code>                                  | Must be either <i>global</i> or <i>local</i> . |

## Additional Options

| Options                                | Description                                                                                                 | Behaviour if omitted                                        |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------|
| <code>--platform</code>                | One of <i>aws</i> , <i>docker</i> , <i>bare</i> .                                                           | Defaults to <i>aws</i> .                                    |
| <code>--location-names</code>          | A space-separated list of location names. The number of locations is equal to the number of names supplied. | TPA will configure a single location with three data nodes. |
| <code>--witness-only-location</code>   | A location name, must be a member of <i>location-names</i> .                                                | No witness-only location is added.                          |
| <code>--data-nodes-per-location</code> | The number of data nodes in each location, must be at least 2.                                              | Defaults to 3.                                              |
| <code>--add-proxy-nodes-per-lo</code>  | The number of proxy nodes in each location.                                                                 | PGD-proxy will be installed on each data node.              |
| <code>--enable-camo</code>             | Sets two data nodes in each location as CAMO partners.                                                      | CAMO will not be enabled.                                   |
| <code>--bdr-database</code>            | The name of the database to be used for replication.                                                        | Defaults to <i>bdrdb</i> .                                  |
| <code>--enable-pgd-probes</code>       | Enable http(s) api endpoints for pgd-proxy such as <i>health/is-ready</i> to allow probing proxy's health.  | Disabled by default.                                        |
| <code>--proxy-listen-port</code>       | The port on which proxy nodes will route traffic to the write leader.                                       | Defaults to 6432                                            |
| <code>--proxy-read-only-port</code>    | The port on which proxy nodes will route read-only traffic to shadow nodes.                                 | Defaults to 6433                                            |

## More detail about PGD-Always-ON configuration

A PGD-Always-ON cluster comprises a number of locations, preferably odd, each with the same number of data nodes, again preferably odd. If you do not specify any `--location-names` , the default is to use a single location with three data nodes.

Location names for the cluster are specified as `--location-names dc1 dc2 ...` . A location represents an independent data centre that provides a level of redundancy, in whatever way this definition makes sense to your use case. For example, AWS regions, your own data centres, or any other designation to identify where your servers are hosted.

Use `--data-nodes-per-location N` to specify the number of data nodes in each location. The minimum number is 2, the default is 3.

If you specify an even number of data nodes per location, TPA will add an extra witness node to each location automatically. This retains the ability to establish reliable consensus while allowing cost savings (a witness has minimal hardware requirements compared to the data nodes).

A cluster with only two locations would entirely lose the ability to establish global consensus if one of the locations were to fail. We recommend adding a third witness-only location (which contains no data nodes, only a witness node, again used to reliably establish consensus). Use `--witness-only-location loc` to designate one of your locations as a witness.

By default, every data node (in every location) will also run PGD-Proxy for connection routing. To create separate

PGD-Proxy instances instead, use `--add-proxy-nodes-per-location 3` (or however many proxies you want to add).

Depending on your use-case, you must specify `--pgd-proxy-routing local` or `global` to configure how PGD-Proxy will route connections to a write leader. Local routing will make every PGD-Proxy route to a write leader within its own location (suitable for geo-sharding applications). Global routing will make every proxy route to a single write leader, elected amongst all available data nodes across all locations.

You may optionally specify `--bdr-database dbname` to set the name of the database with BDR enabled (default: `bdrdb`).

You may optionally specify `--enable-camo` to set two data nodes in each region as CAMO partners.

You may optionally specify `--enable-pgd-probes [{http, https}]` to enable http(s) api endpoints that will allow to easily probe proxy's health.

You may also specify any of the options described by *Cluster configuration*.

## 3.2 PGD Lightweight

### Note

This architecture is for Postgres Distributed 5 only. If you require PGD 4 or 3.7 please use *BDR-Always-ON*.

EDB Postgres Distributed 5 in a Lightweight configuration, suitable for use in test and production.

This architecture requires an EDB subscription. All software will be sourced from *Configuring EDB Repos 2.0 repositories*.

### 3.2.1 Cluster configuration

#### Overview of configuration options

An example invocation of `tpaexec configure` for this architecture is shown below.

```
tpaexec configure ~/clusters/pgd-lw \
  --architecture Lightweight \
  --edb-postgres-extended 15 \
  --platform aws --instance-type t3.micro \
  --distribution Debian \
  --location-names main dr \
```

You can list all available options using the help command.

```
tpaexec configure --architecture Lightweight --help
```

The table below describes the mandatory options for PGD-Always-ON and additional important options. More detail on the options is provided in the following section.

#### Mandatory Options

| Options                          | Description                       |
|----------------------------------|-----------------------------------|
| <code>--architecture (-a)</code> | Must be set to <i>Lightweight</i> |

continues on next page

Table 3 – continued from previous page

| Options                                                           | Description                                   |
|-------------------------------------------------------------------|-----------------------------------------------|
| Postgres flavour and version (e.g. <code>--postgresql 15</code> ) | A valid <i>Postgres flavour and version</i> . |

### Additional Options

| Options                               | Description                                                                                                 | Behaviour if omitted                                        |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------|
| <code>--platform</code>               | One of <i>aws</i> , <i>docker</i> , <i>bare</i> .                                                           | Defaults to <i>aws</i> .                                    |
| <code>--location-names</code>         | A space-separated list of location names. The number of locations is equal to the number of names supplied. | TPA will configure a single location with three data nodes. |
| <code>--add-proxy-nodes-per-lo</code> | The number of proxy nodes in each location.                                                                 | PGD-proxy will be installed on each data node.              |
| <code>--bdr-database</code>           | The name of the database to be used for replication.                                                        | Defaults to <i>bdrdb</i> .                                  |
| <code>--enable-pgd-probes</code>      | Enable http(s) api endpoints for pgd-proxy such as <i>health/is-ready</i> to allow probing proxy's health.  | Disabled by default.                                        |
| <code>--proxy-listen-port</code>      | The port on which proxy nodes will route traffic to the write leader.                                       | Defaults to 6432                                            |
| <code>--proxy-read-only-port</code>   | The port on which proxy nodes will route read-only traffic to shadow nodes.                                 | Defaults to 6433                                            |

### More detail about Lightweight configuration

A PGD Lightweight cluster comprises 2 locations, with a primary active location containing 2 nodes and a disaster recovery (dr) location with a single node.

Location names for the cluster are specified as `--location-names primary dr` . A location represents an independent data centre that provides a level of redundancy, in whatever way this definition makes sense to your use case. For example, AWS regions, your own data centres, or any other designation to identify where your servers are hosted.

By default, every data node (in every location) will also run PGD-Proxy for connection routing. To create separate PGD-Proxy instances instead, use `--add-proxy-nodes-per-location 3` (or however many proxies you want to add).

Global routing will make every proxy route to a single write leader, elected amongst all available data nodes across all locations.

You may optionally specify `--bdr-database dbname` to set the name of the database with BDR enabled (default: *bdrdb*).

You may optionally specify `--enable-pgd-probes [{http, https}]` to enable http(s) api endpoints that will allow to easily probe proxy's health.

You may also specify any of the options described by *Cluster configuration* .

## 3.3 BDR-Always-ON

EDB Postgres Distributed 3.7 or 4 in an Always-ON configuration, suitable for use in test and production.

This architecture requires an EDB subscription. All software will be sourced from *Configuring EDB Repos 2.0 repositories* .

The BDR-Always-ON architecture has four variants, which can be selected with the `--layout` configure option:

1. bronze: 2×bdr+primary, bdr+witness, barman, 2×harp-proxy

2. silver: bronze, with bdr+witness promoted to bdr+primary, and barman moved to separate location
3. gold: two symmetric locations with 2×bdr+primary, 2×harp-proxy, and barman each; plus a bdr+witness in a third location
4. platinum: gold, but with one bdr+readonly (logical standby) added to each of the main locations

You can check EDB's Postgres Distributed Always On Architectures

[whitepaper](#)

for the detailed layout diagrams.

This architecture is meant for use with PGD versions 3.7 and 4.

### 3.3.1 Cluster configuration

#### Overview of configuration options

An example invocation of `tpaexec configure` for this architecture is shown below.

```
tpaexec configure ~/clusters/bdr \
  --architecture BDR-Always-ON \
  --platform aws --region eu-west-1 --instance-type t3.micro \
  --distribution Debian \
  --edb-postgres-advanced 14 --redwood
  --layout gold \
  --harp-consensus-protocol bdr
```

You can list all available options using the help command.

```
tpaexec configure --architecture BDR-Always-ON --help
```

The tables below describe the mandatory options for BDR-Always-ON and additional important options. More detail on the options is provided in the following section.

#### Mandatory Options

| Option                                                            | Description                                                            |
|-------------------------------------------------------------------|------------------------------------------------------------------------|
| <code>--architecture (-a)</code>                                  | Must be set to <i>BDR-Always-ON</i> .                                  |
| Postgres flavour and version (e.g. <code>--postgresql 14</code> ) | A valid <i>Postgres flavour and version</i> .                          |
| <code>--layout</code>                                             | One of <i>bronze</i> , <i>silver</i> , <i>gold</i> , <i>platinum</i> . |
| <code>--harp-consensus-protocol</code>                            | One of <i>bdr</i> , <i>etcd</i> .                                      |

#### Additional Options

| Option                      | Description                                            | Behaviour if omitted       |
|-----------------------------|--------------------------------------------------------|----------------------------|
| <code>--platform</code>     | One of <i>aws</i> , <i>docker</i> , <i>bare</i> .      | Defaults to <i>aws</i> .   |
| <code>--enable-camo</code>  | Sets two data nodes in each location as CAMO partners. | CAMO will not be enabled.  |
| <code>--bdr-database</code> | The name of the database to be used for replication.   | Defaults to <i>bdrdb</i> . |

continues on next page

Table 6 – continued from previous page

| Option                            | Description                                                                                          | Behaviour if omitted |
|-----------------------------------|------------------------------------------------------------------------------------------------------|----------------------|
| <code>--enable-harp-probes</code> | Enable http(s) api endpoints for harp such as <i>health/is-ready</i> to allow probing harp's health. | Disabled by default. |

### More detail about BDR-Always-ON configuration

You must specify `--layout layoutname` to set one of the supported BDR use-case variations. The permitted arguments are bronze, silver, gold, and platinum. The bronze, gold and platinum layouts have a PGD witness node to ensure odd number of nodes for Raft consensus majority. Witness nodes do not participate in the data replication.

You must specify `--harp-consensus-protocol protocolname`. The supported protocols are bdr and etcd; see [Configuring HARP](#) for more details.

You may optionally specify `--bdr-database dbname` to set the name of the database with PGD enabled (default: bdrdb).

You may optionally specify `--enable-camo` to set the pair of PGD primary instances in each region to be each other's CAMO partners.

You may optionally specify `--enable-harp-probes [{http, https}]` to enable http(s) api endpoints that will allow to easily probe harp's health.

Please note we enable HARP2 by default in BDR-Always-ON architecture.

You may also specify any of the options described by [Cluster configuration](#).

## 3.4 M1

A Postgres cluster with a single primary node and physical replication to a number of standby nodes including backup and failover management.

This architecture is suitable for production and is also suited to testing, demonstrating and learning due to its simplicity and ability to be configured with no proprietary components.

If you select subscription-only EDB software with this architecture it will be sourced from EDB Repos 2.0 and you will need to [Configure EDB Repos 2.0 repositories](#).

### 3.4.1 Failover management

The M1 architecture always includes a failover manager. Supported options are repmgr, EDB Failover Manager (EFM) and Patroni. In all cases, the failover manager will be configured by default to ensure that a replica will be promoted to take the place of the primary should the primary become unavailable.

#### Application failover

The M1 architecture does not generally provide an automatic facility to reroute application traffic to the primary. There are several ways you can add this capability to your cluster.

In TPA:

- If you choose repmgr as the failover manager and enable PgBouncer, you can include the `repmgr_redirect_pgouncer: true` hash under `cluster_vars` in `config.yml`. This causes repmgr to automatically reconfigure PgBouncer to route traffic to the new primary on failover.
- If you choose Patroni as the failover manager and enable PgBouncer, Patroni will automatically reconfigure PgBouncer to route traffic to the new primary on failover.

- If you choose EFM as the failover manager, you can use the `efm_conf_settings` hash under `cluster_vars` in `config.yml` to *configure EFM to use a virtual IP address (VIP)* . This is an additional IP address which will always route to the primary node.
- Place an appropriate proxy or load balancer between the cluster and you application and use a *TPA hooks* to configure your selected failover manager to update it with the route to the new primary on failover.
- Handle failover at the application itself, for example by using multi-host connection strings.

### Backup failover

TPA does not configure any kind of ‘backup failover’. If the Postgres node from which you are backing up is down, backups will simply halt until the node is back online. To manually connect the backup to the new primary, edit `config.yml` to add the `backup` hash to the new primary instance and re-run `tpaexec deploy`.

## 3.4.2 Cluster configuration

### Overview of configuration options

An example invocation of `tpaexec configure` for this architecture is shown below.

```
tpaexec configure ~/clusters/m1 \
  --architecture M1 \
  --platform aws --region eu-west-1 --instance-type t3.micro \
  --distribution Debian \
  --postgresql 14 \
  --failover-manager repmgr \
  --data-nodes-per-location 3
```

You can list all available options using the help command.

```
tpaexec configure --architecture M1 --help
```

The tables below describe the mandatory options for M1 and additional important options. More detail on the options is provided in the following section.

### Mandatory Options

| Parameter                                                         | Description                                                                                                                                 |
|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <code>--architecture (-a)</code>                                  | Must be set to <i>M1</i> .                                                                                                                  |
| Postgres flavour and version (e.g. <code>--postgresql 15</code> ) | A valid <i>Postgres flavour and version</i> .                                                                                               |
| One of: <code>--failover-manager {efm, repmgr, patroni}</code>    | Select the failover manager from <i>Configuring EFM</i> , <i>Repmgr redirect pgbouncer</i> and <i>Patroni cluster management commands</i> . |
| <code>--enable-efm</code>                                         |                                                                                                                                             |
| <code>--enable-repmgr</code>                                      |                                                                                                                                             |
| <code>--enable-patroni</code>                                     |                                                                                                                                             |

### Additional Options

| Parameter                              | Description                                                                                                                                                                            | Behaviour if omitted                             |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------|
| <code>--platform</code>                | One of <i>aws</i> , <i>docker</i> , <i>bare</i> .                                                                                                                                      | Defaults to <i>aws</i> .                         |
| <code>--location-names</code>          | A space-separated list of location names. The number of locations is equal to the number of names supplied.                                                                            | A single location called "main" is used.         |
| <code>--primary-location</code>        | The location where the primary server will be. Must be a member of <i>location-names</i> .                                                                                             | The first listed location is used.               |
| <code>--data-nodes-per-location</code> | A number from 1 upwards. In each location, one node will be configured to stream directly from the cluster's primary node, and the other nodes, if present, will stream from that one. | Defaults to 2.                                   |
| <code>--witness-only-location</code>   | A location name, must be a member of <i>location-names</i> . This location will be populated with a single witness node only.                                                          | No witness-only location is added.               |
| <code>--single-node-location</code>    | A location name, must be a member of <i>location-names</i> . This location will be populated with a single data node only.                                                             | No single-node location is added.                |
| <code>--enable-haproxy</code>          | Two additional nodes will be added as a load balancer layer. Only supported with Patroni as the failover manager.                                                                      | HAproxy nodes will not be added to the cluster.  |
| <code>--enable-pgbouncer</code>        | PgBouncer will be configured in the Postgres nodes to pool connections for the primary.                                                                                                | PgBouncer will not be configured in the cluster. |
| <code>--patroni-dcs</code>             | Select the Distributed Configuration Store backend for patroni. Only option is <i>etcd</i> at this time. Only supported with Patroni as the failover manager.                          | Defaults to <i>etcd</i> .                        |
| <code>--efm-bind-by-hostname</code>    | Enable efm to use hostnames instead of IP addresses to configure the cluster <i>bind.address</i> .                                                                                     | Defaults to use IP addresses                     |

### More detail about M1 configuration

You may also specify any of the options described by [Cluster configuration](#).

## 3.5 PGD-S

!!!Note

This architecture is for Postgres Distributed 6 only. If you require PGD 5 please use [PGD-Always-ON](#).

This architecture requires an EDB subscription. All software is sourced from [Configuring EDB Repos 2.0 repositories](#).

### 3.5.1 Cluster configuration

#### Overview of configuration options

An example invocation of `tpaexec configure` for this architecture is shown below.

```
tpaexec configure ~/clusters/pgd-s \
  --architecture PGD-S
  --edb-postgres-extended 15 \
  --platform aws --instance-type t3.micro \
  --distribution Debian \
```

You can list all available options using the help command.

```
tpaexec configure --architecture PGD-S --help
```

The table below describes the mandatory options for PGD-S and additional important options. More detail on the options is provided in the following section.

### Mandatory Options

| Options                                                           | Description                                   |
|-------------------------------------------------------------------|-----------------------------------------------|
| <code>--architecture (-a)</code>                                  | Must be set to <i>PGD-S</i>                   |
| Postgres flavour and version (e.g. <code>--postgresql 15</code> ) | A valid <i>Postgres flavour and version</i> . |

### Additional Options

| Options                                  | Description                                                              | Behaviour if omitted                                                   |
|------------------------------------------|--------------------------------------------------------------------------|------------------------------------------------------------------------|
| <code>--platform</code>                  | One of <i>aws</i> , <i>docker</i> , <i>bare</i> .                        | Defaults to <i>aws</i> .                                               |
| <code>--bdr-database</code>              | The name of the database to be used for replication.                     | Defaults to <i>bdrdb</i> .                                             |
| <code>--layout</code>                    | <i>standard</i> or <i>near-far</i>                                       | Defaults to <i>standard</i>                                            |
| <code>--add-subscriber-only-nodes</code> | The number of subscriber-only nodes to add, up to 10.                    | Defaults to 0                                                          |
| <code>--read-write-port</code>           | The port for Connection Manager to listen on for read-write connections. | Left empty in config.yml, allowing default of the postgres port + 1000 |
| <code>--read-only-port</code>            | The port for Connection Manager to listen on for read-only connections.  | Left empty in config.yml, allowing default of the read-write port + 1  |
| <code>--http-port</code>                 | The port for Connection Manager to listen on for http api connections.   | Left empty in config.yml, allowing default of the read-write port + 2  |
| <code>--use https</code>                 | Enable https for Connection Manager's http api                           | https is not enabled                                                   |

### More detail about PGD-S configuration

A PGD-S cluster has three data nodes. In the *standard* layout the nodes are all in the same location; in the *near-far* layout two nodes are in the primary location and the other is in a secondary location. See the [PGD documentation](#) for more information about the two layouts.

The cluster also contains one barman node and up to 10 subscriber-only nodes, controlled by the `--add-subscriber-only-nodes` parameter. These are always in the primary location.

You may optionally specify `--bdr-database dbname` to set the name of the database with BDR enabled (default: *bdrdb*).

You may also specify any of the options described by [Cluster configuration](#) .

## 3.6 PGD-X

!!!Note

This architecture is for Postgres Distributed 6 only. If you require PGD 5 please use *PGD-Always-ON* .

This architecture requires an EDB subscription. All software is sourced from [Configuring EDB Repos 2.0 repositories](#) .

### 3.6.1 Cluster configuration

#### Overview of configuration options

An example invocation of `tpaexec configure` for this architecture is shown below.

```
tpaexec configure ~/clusters/pgd-x \
  --architecture PGD-X \
  --edb-postgres-extended 15 \
  --platform aws --instance-type t3.micro \
  --distribution Debian \
  --pgd-routing global \
  --location-names dc1 dc2 dc3 \
  --witness-only-location dc3 \
  --data-nodes-per-location 2
```

You can list all available options using the help command.

```
tpaexec configure --architecture PGD-X --help
```

The table below describes the mandatory options for PGD-X and additional important options. More detail on the options is provided in the following section.

#### Mandatory Options

| Options                                                           | Description                                    |
|-------------------------------------------------------------------|------------------------------------------------|
| <code>--architecture (-a)</code>                                  | Must be set to <i>PGD-X</i>                    |
| Postgres flavour and version (e.g. <code>--postgresql 15</code> ) | A valid <i>Postgres flavour and version</i> .  |
| <code>--pgd-routing</code>                                        | Must be either <i>global</i> or <i>local</i> . |

#### Additional Options

| Options                                | Description                                                                                                 | Behaviour if omitted                                        |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------|
| <code>--platform</code>                | One of <i>aws</i> , <i>docker</i> , <i>bare</i> .                                                           | Defaults to <i>aws</i> .                                    |
| <code>--location-names</code>          | A space-separated list of location names. The number of locations is equal to the number of names supplied. | TPA will configure a single location with three data nodes. |
| <code>--witness-only-location</code>   | A location name, must be a member of <i>location-names</i> .                                                | No witness-only location is added.                          |
| <code>--data-nodes-per-location</code> | The number of data nodes in each location, must be at least 2.                                              | Defaults to 3.                                              |
| <code>--enable-camo</code>             | Sets two data nodes in each location as CAMO partners.                                                      | CAMO will not be enabled.                                   |
| <code>--bdr-database</code>            | The name of the database to be used for replication.                                                        | Defaults to <i>bdrdb</i> .                                  |
| <code>--enable-pgd-probes</code>       | Enable http(s) api endpoints for pgd-proxy such as <i>health/is-ready</i> to allow probing proxy's health.  | Disabled by default.                                        |
| <code>--proxy-listen-port</code>       | The port on which proxy nodes will route traffic to the write leader.                                       | Defaults to 6432                                            |
| <code>--proxy-read-only-port</code>    | The port on which proxy nodes will route read-only traffic to shadow nodes.                                 | Defaults to 6433                                            |

### More detail about PGD-X configuration

A PGD-X cluster comprises a number of locations, preferably odd, each with the same number of data nodes, again preferably odd. If you do not specify any `--location-names`, the default is to use a single location with three data nodes.

Location names for the cluster are specified as `--location-names dc1 dc2 ...`. A location represents an independent data center that provides a level of redundancy, in whatever way this definition makes sense to your use case. For example, AWS regions, your own data centers, or any other designation to identify where your servers are hosted.

Use `--data-nodes-per-location N` to specify the number of data nodes in each location. The minimum number is 2, the default is 3.

If you specify an even number of data nodes per location, TPA will add an extra witness node to each location automatically. This retains the ability to establish reliable consensus while allowing cost savings (a witness has minimal hardware requirements compared to the data nodes).

A cluster with only two locations would entirely lose the ability to establish global consensus if one of the locations were to fail. We recommend adding a third witness-only location (which contains no data nodes, only a witness node, again used to reliably establish consensus). Use `--witness-only-location loc` to designate one of your locations as a witness.

Depending on your use-case, you must specify `--pgd-routing local` or `global` to configure how Connection Manager will route connections to a write leader. Local routing will make every Connection Manager route to a write leader within its own location (suitable for geo-sharding applications). Global routing will make every Connection Manager route to a single write leader, elected amongst all available data nodes across all locations.

You may optionally specify `--bdr-database dbname` to set the name of the database with BDR enabled (default: `bdrdb`).

You may optionally specify `--enable-camo` to set two data nodes in each region as CAMO partners.

You may optionally specify `--enable-pgd-probes [{http, https}]` to enable http(s) api endpoints that will allow to easily probe proxy's health.

You may also specify any of the options described by *[Cluster configuration](#)*.



### 4.1 aws

TPA fully supports provisioning production clusters on AWS EC2.

#### 4.1.1 API access setup

To use the AWS API, you must:

- Obtain an access keypair
- Add it to your configuration

For example,

```
cat > ~/.aws/credentials
```

```
[default]
aws_access_key_id = AKIAIOSFODNN7EXAMPLE
aws_secret_access_key = wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
```

The IAM user should at least have following set of permissions so tpaexec can use it to provision ec2 resources.

```
ec2:AssociateRouteTable
ec2:AttachInternetGateway
ec2:AuthorizeSecurityGroupIngress
ec2:CreateInternetGateway
ec2:CreateRoute
ec2:CreateRouteTable
ec2:CreateSecurityGroup
ec2:CreateSubnet
ec2:CreateTags
ec2:CreateVpc
ec2>DeleteKeyPair
```

(continues on next page)

(continued from previous page)

```
ec2:DeleteRouteTable
ec2:DeleteSecurityGroup
ec2:DeleteSubnet
ec2:DeleteVpc
ec2:DescribeImages
ec2:DescribeInstanceStatus
ec2:DescribeInstances
ec2:DescribeInternetGateways
ec2:DescribeKeyPairs
ec2:DescribeRouteTables
ec2:DescribeSecurityGroups
ec2:DescribeSubnets
ec2:DescribeTags
ec2:DescribeVolumes
ec2:DescribeVpcAttribute
ec2:DescribeVpcClassicLink
ec2:DescribeVpcClassicLinkDnsSupport
ec2:DescribeVpcs
ec2:DisassociateRouteTable
ec2:ImportKeyPair
ec2:ModifyVpcAttribute
ec2:RevokeSecurityGroupIngress
ec2:RunInstances
ec2:TerminateInstances
iam:AddRoleToInstanceProfile
iam:CreateInstanceProfile
iam:CreateRole
iam>DeleteInstanceProfile
iam>DeleteRole
iam>DeleteRolePolicy
iam:GetInstanceProfile
iam:GetRole
iam:GetRolePolicy
iam:ListAttachedRolePolicies
iam:ListGroups
iam:ListInstanceProfiles
iam:ListInstanceProfilesForRole
iam:ListRolePolicies
iam:ListRoles
iam:ListUsers
iam:PassRole
iam:PutRolePolicy
iam:RemoveRoleFromInstanceProfile
kms:CreateGrant
kms:GenerateDataKeyWithoutPlaintext
s3:CreateBucket
s3:GetBucketVersioning
s3:GetObject
s3:GetObjectTagging
s3:ListAllMyBuckets
s3:ListBucket
s3:ListBucketVersions
```

(continues on next page)

(continued from previous page)

```
s3:PutBucketOwnershipControls
s3:PutObject
s3:PutObjectAcl
```

### 4.1.2 Introduction

The service is physically subdivided into [regions and availability zones](#) . An availability zone is represented by a region code followed by a single letter, e.g., eu-west-1a (but that name may refer to different locations for different AWS accounts, and there is no way to coordinate the interpretation between accounts).

AWS regions are completely isolated from each other and share no resources. Availability zones within a region are physically separated, and logically mostly isolated, but are connected by low-latency links and are able to share certain networking resources.

#### Networking

All networking configuration in AWS happens in the context of a [Virtual Private Cloud](#)

within a region. Within a VPC, you can create [subnets](#)

that is tied to a specific availability zone, along with internet gateways, routing tables, and so on.

You can create any number of [Security Groups](#)

to configure rules for what inbound and outbound traffic is permitted to instances (in terms of protocol, a destination port range, and a source or destination IP address range).

#### Instances

AWS EC2 offers a variety of [instance types](#)

with different hardware configurations at different price/performance points. Within a subnet in a particular availability zone, you can create [EC2 instances](#)

based on a distribution image known as an [AMI](#) , and attach one or more [EBS volumes](#)

to provide persistent storage to the instance. You can SSH to the instances by registering an [SSH public key](#) .

Instances are always assigned a private IP address within their subnet. Depending on the subnet configuration, they may also be assigned an [ephemeral public IP address](#)

(which is lost when the instance is shut down, and a different ephemeral IP is assigned when it is started again). You can instead assign a static region-specific routable IP address known as an [Elastic IP](#)

to any instance.

For an instance to be reachable from the outside world, it must not only have a routable IP address, but the VPC's networking configuration (internet gateway, routing tables, security groups) must also align to permit access.

### 4.1.3 Configuration

Here's a brief description of the AWS-specific settings that you can specify via `tpaexec configure` or define directly in `config.yml`.

## Regions

You can specify one or more regions for the cluster to use with `--region` or `--regions`. TPA will generate the required vpc entries associated to each of them and distribute locations into these regions evenly by using different availability zones while possible.

`regions` are different from `locations`, each location belongs to a region (and an availability zone inside this region). `regions` are AWS specific objects, `locations` are cluster objects.

Note: When specifying multiple regions, you need to manually edit network configurations:

- `ec2_vpc` entries must have non-overlapping cidr networks to allow use of AWS vpc peering. by default TPA will set all cidr to `10.33.0.0/16`. See [VPC \(required\)](#) for more informations.
- each `location` must be updated with `subnet` that match the `ec2_vpc cidr` they belong to. See [Subnets \(optional\)](#) for more informations.
- TPA creates security groups with basic rules under `cluster_rules` and those need to be updated to match `ec2_vpc cidr` for each `subnet cidr`. see [Security groups \(optional\)](#) for more informations.
- VPC peering must be setup manually before `tpaexec deploy`. We recommend creating VPCs and required VPC peerings before running `tpaexec configure` and using `vpc-id` in `config.yml`. See [VPC \(required\)](#) for more informations.

## VPC (required)

You must specify a VPC to use:

```
ec2_vpc:
  Name: Test
  cidr: 10.33.0.0/16
```

This is the default configuration, which creates a VPC named Test with the given CIDR if it does not exist, or uses the existing VPC otherwise.

To create a VPC, you must specify both the Name and the cidr. If you specify only a VPC Name, TPA will fail if a matching VPC does not exist.

If TPA creates a VPC, `tpaexec deprovision` will attempt to remove it, but will leave any pre-existing VPC alone. (Think twice before creating new VPCs, because AWS has a single-digit default limit on the number of VPCs per account.)

If you need more fine-grained matching, or to specify different VPCs in different regions, you can use the expanded form:

```
ec2_vpc:
  eu-west-1:
    Name: Test
    cidr: 172.16.0.0/16
  us-east-1:
    filters:
      vpc-id: vpc-nnn
  us-east-2:
    Name: Example
    filters:
      [filter expressions]
```

### AMI (required)

You must specify an AMI to use:

```
ec2_ami:
  Name: xxx
  Owner: self
```

You can add filter specifications for more precise matching:

```
ec2_ami:
  Name: xxx
  Owner: self
  filters:
    architecture: x86_64
    [more key/value filters]
```

(By default, `tpaexec configure` will select a suitable `ec2_ami` for you based on the `--distribution` argument.)

### Subnets (optional)

Every instance must specify its subnet (in CIDR form, or as a subnet-xxx id). You may optionally specify the name and availability zone for each subnet that we create:

```
ec2_vpc_subnets:
  us-east-1:
    192.0.2.0/27:
      az: us-east-1b
      Name: example1
    192.0.2.100/27:
      az: us-east-1b
      Name: example2
```

### Security groups (optional)

By default, we create a security group for the cluster. To use one or more existing security groups, set:

```
ec2_groups:
  us-east-1:
    group-name:
      - foo
```

If you want to customise the rules in the default security group, set `cluster_rules`:

```
cluster_rules:
- cidr_ip: 0.0.0.0/0
  from_port: 22
  proto: tcp
  to_port: 22
- cidr_ip: 192.0.2.0/27
  from_port: 0
  proto: tcp
  to_port: 65535
```

(continues on next page)

(continued from previous page)

```
- cidr_ip: 192.0.2.100/27
  from_port: 0
  proto: tcp
  to_port: 65535
```

This example permits ssh (port 22) from any address, and TCP connections on any port from specific IP ranges. (Note: from\_port and to\_port define a numeric range of ports, not a source and destination.)

If you set up custom rules or use existing security groups, you must ensure that instances in the cluster are allowed to communicate with each other as required (e.g., allow tcp/5432 for Postgres).

### Internet gateways (optional)

By default, we create internet gateways for every VPC, unless you set:

```
ec2_instance_reachability: private
```

For more fine-grained control, you can set:

```
ec2_vpc_igw:
  eu-west-1: yes
  eu-central-1: yes
  us-east-1: no
```

### SSH keys (optional)

```
# Set this to change the name under which we register our SSH key.

# ec2_key_name: tpa_cluster_name

#
# Set this to use an already-registered key.

# ec2_instance_key: xxx
```

### S3 bucket (optional)

TPA requires access to an S3 bucket to provision an AWS cluster. This bucket is used to temporarily store files such as SSH host keys, but may also be used for other cluster data (such as backups).

By default, TPA will use an S3 bucket named `edb-tpa-<aws-account-user-id>` for any clusters you provision. (If the bucket does not exist, you will be asked to confirm that you want TPA to create it for you.)

To use an existing S3 bucket instead, set

```
cluster_bucket: name-of-bucket
```

(You can also set `cluster_bucket: auto` to accept the default bucket name without the confirmation prompt.)

TPA will never remove any S3 buckets when you deprovision the cluster. To remove the bucket yourself, run:

```
aws s3 rb s3://<bucket> --force
```

The IAM user you are using to provision the instances must have read and write access to this bucket. During provisioning, tpaexec will provide instances with read-only access to the cluster\_bucket through the instance profile.

### Elastic IP addresses

To use elastic IP addresses, set `assign_elastic_ip` to `true` in `config.yml`, either in `instance_defaults` to affect all the instances in your cluster or individually on the separate instances as required. By default, this will allocate a new elastic ip address and assign it to the new instance. To use an elastic IP address that has already been allocated but not yet assigned, use `elastic_ip: 34.252.55.252`, substituting in your allocated address.

### Instance profile (optional)

```
# Set this to change the name of the instance profile role we create.

# cluster_profile: cluster_name_profile

#

# Set this to use an existing instance profile (which must have all the
# required permissions assigned to it).

# instance_profile_name: xxx
```

## 4.2 bare(-metal servers)

Set `platform: bare` in `config.yml`

This platform is meant to support any server that is accessible via SSH, including bare-metal servers as well as already-provisioned servers on any cloud platform (including AWS).

You must define the IP address(es) and username for each target server:

```
instances:
- node: 1
  Name: igor
  platform: bare
  public_ip: 192.0.2.1
  private_ip: 192.0.2.222
  vars:
    ansible_user: xyzzy
```

You must ensure that

1. TPA can ssh to the instance as `ansible_user`
2. The `ansible_user` has sudo access on the instance

### 4.2.1 SSH access

In the example above, TPA will ssh to `xyzzy@192.0.2.1` to access the instance.

By default, TPA will run `ssh-keygen` to generate a new SSH keypair in your cluster directory. The private key is named `id_cluster_name` and the public key is stored in `id_cluster_name.pub`.

You must either set `ssh_key_file: /path/to/id_keyname` to use a different key that the instance will accept, or configure the instance to allow access from the generated key (e.g., use `ssh-copy-id`, which will append the contents of `id_cluster_name.pub` to `~xyzzy/.ssh/authorized_keys`).

You must also ensure that `ssh` can verify the host key(s) of the instance. You can either add entries to the `known_hosts` file in your cluster directory, or install the TPA-generated host keys from `hostkeys/ssh_host_*_key*` in your cluster directory into `/etc/ssh` on the instance (the generated `tpa_known_hosts` file contains entries for these keys).

For example, to `ssh` in with the generated user key, but keep the existing host keys, you can do:

```
cd ~/clusters/speedy
ssh-copy-id -i id_speedy xyzzy@192.0.2.1
ssh-keyscan -H 192.0.2.1 >> tpa_known_hosts
```

Run `tpaexec ping ~/clusters/speedy` to check if it's working. If not, append `-vvv` to the command to look at the complete `ssh` command-line. (Note: Ansible will invoke `ssh` to execute a command like `bash -c 'python3 && sleep 0'` on the instance. If you run `ssh` commands by hand while debugging, replace this with a command that produces some output and then exits instead, e.g., `'id'`.)

For more details:

- *ssh\_key\_file*
- *Managing SSH host keys*

## 4.2.2 Distribution support

TPA will try to detect the distribution running on target instances, and fail if it is not supported. TPA currently supports Debian (10/11/12; or buster/bullseye/bookworm), Ubuntu (16.04/18.04/20.04/22.04; or xenial/bionic/focal/jammy), and RHEL/CentOS/Rocky/AlmaLinux (7.x/8.x) on bare instances.

### 4.2.3 IP addresses

You can specify the `public_ip`, `private_ip`, or both for any instance.

TPA uses these IP addresses in two ways: first, to `ssh` to the instance to execute commands during deployment; and second, to set up communications within the cluster, e.g., for `/etc/hosts` or to set `primary_conninfo` for Postgres.

If you specify a `public_ip`, it will be used to `ssh` to the instances during deployment. If you specify a `private_ip`, it will be used to set up communications within the cluster. If you specify both, the `public_ip` will be used during deployment, and the `private_ip` for cluster communications.

If you specify only one or the other, the address will be used for both purposes. For example, you could set only `public_ip` for servers on different networks, or only `private_ip` if you're running TPA inside a closed network. (Instead of using `public/private`, you can set `ip_address` if you need to specify only one IP address.)

### 4.2.4 Starting afresh

To start afresh with a cluster on the bare platform, use the appropriate external tools to reinstall, reimage, or reprovision the servers, and repeat the process described in this document. If your new servers have different IP addresses or if you have a complex `ssh` setup, it may be easier to run *tpaexec deprovision* to remove all the locally created files and then *tpaexec provision*

to recreate them, followed by repeating the process from this document, as above.

## 4.3 Docker

TPA can create Docker containers and deploy a cluster to them. At present, it sets up containers to run systemd and other services as if they were ordinary VMs.

Deploying to docker containers is an easy way to test different cluster configurations. It is not meant for production use.

### 4.3.1 Synopsis

Just select the platform at configure-time:

```
tpaexec configure clustername --platform docker [...]
tpaexec provision clustername
tpaexec deploy clustername
```

### 4.3.2 Operating system selection

Use the standard `--os Debian/Ubuntu/RedHat/SLES` configure option to select which distribution to use for the containers. TPA will build its own systemd-enabled images for this distribution. These images will be named with a `tpa/` prefix, e.g., `tpa/redhat:8`.

Use `--os-image some/image:name` to specify an existing systemd-enabled image instead. For example, the `centos/systemd`

image (based on CentOS 7) can be used in this way.

### 4.3.3 Installing Docker

We test TPA with the latest stable Docker-CE packages.

This documentation assumes that you have a working Docker installation, and are familiar with basic operations such as pulling images and creating containers.

Please consult the [Docker documentation](#) if you need help to [install Docker](#) and [get started](#) with it.

On MacOS X, you can [install "Docker Desktop forMac"](#)

and launch Docker from the application menu.

### Cgroups

TPA supports Docker containers on hosts running cgroups version 1 or 2. On a host running cgroups2, instances running RHEL 7 are not supported.

If you need to use RHEL 7 instances but your host is running cgroups version 2, you can switch to cgroups version 1 as follows.

On Debian-family Linux distributions:

```
echo GRUB_CMDLINE_LINUX=systemd.unified_cgroup_hierarchy=false > \
/etc/default/grub.d/cgroup.cfg
update-grub
reboot
```

On RedHat-family Linux distributions:

```
grubby --args=systemd.unified_cgroup_hierarchy=false --update-kernel=ALL
reboot
```

On MacOS:

1. Edit `~/Library/Group\ Containers/group.com.docker/settings.json` and make the following replacement `"deprecatedCgroupv1": false → "deprecatedCgroupv1": true`
2. Restart Docker Desktop app

## Permissions

TPA expects the user running it to have permission to access to the Docker daemon (typically by being a member of the `docker` group that owns `/var/run/docker.sock`). Run a command like this to check if you have access:

```
docker version --format {{.Server.Version}}
```

```
19. 03.12
```

### !!!Warning

Giving a user the ability to speak to the Docker daemon lets them trivially gain root on the Docker host. Only trusted users should have access to the Docker daemon.

## Privileged containers

By default TPA provisions Docker containers in unprivileged mode, with no added Linux capabilities flags. Such containers cannot manage host firewall rules, file systems, block devices, or most other tasks that require true root privileges on the host.

If you require your containers to run in privileged mode, set the `privileged` boolean variable for the instance(s) that need it, or globally in `instance_defaults`, e.g.:

```
instance_defaults:
  privileged: true
```

### !!!Warning

Running containers in privileged mode allows the root user or any process that can gain root to load kernel modules, modify host firewall rules, escape the container namespace, or otherwise act much as the real host “root” user would. Do not run containers in privileged mode unless you really need to.

## security\_opts and the no-new-privileges flag

`tpaexec` can start docker containers in a restricted mode where processes cannot increase their privileges. `setuid` binaries are restricted, etc. Enable this in `tpaexec` with the `instance_defaults` or per-container variable `docker_security_opts`:

```
instance_defaults:
  docker_security_opts:
    - no-new-privileges
```

Other arguments to `docker run`’s `--security-opts` are also accepted, e.g. SELinux user and role.

## Linux capabilities flags

`tpaexec` exposes Docker’s control over Linux capabilities flags with the `docker_cap_add` list variable, which may be set per-container or in `instance_defaults`. See `man capabilities`, the `docker run` documentation and the documentation for the Ansible `docker_containers` module for details on capabilities flags.

Docker’s `--cap-drop` is also supported via the `docker_cap_drop` list.

For example, to run a container as unprivileged, but give it the ability to modify the system clock, you might write:

```
instance_defaults:
  privileged: false
  docker_cap_add:
    - sys_time
  docker_cap_drop:
    - all
```

## Docker storage configuration

The default Docker configuration on many hosts uses `lvm-loop` block storage and is not suitable for production deployments. Run `docker info` to check which storage driver you are using. If you are using the loopback scheme, you will see something like this:

```
Storage Driver: devicemapper
...
Data file: /dev/loop0
```

Consult the Docker documentation for more information on storage configuration:

- [Storage Drivers](#)
- [Configuring lvm-direct for production](#)

## Docker MTU settings

By default, Docker networks have a Maximum Transmission Unit (MTU) of 1500 bytes. If this is greater than the MTU of your host system's network interface you may experience problems routing connections through that interface to Docker containers. You can check the MTU of your network interfaces using the command `ipconfig | grep mtu, ip | grep mtu` or similar. You can change the MTU of a Docker network provisioned by TPA by adding the appropriate driver options to the network in `config.yml` as shown below.

```
docker_networks:
- ipam_config:
  - subnet: 10.33.214.192/28
  name: tpa-docker
  driver_options:
    com.docker.network.driver.mtu: 1400
```

!!!Warning

The MTU can only be set when the Docker network is first provisioned. Subsequent changes in `config.yml` will have no effect.

## 4.3.4 Docker container management

All of the docker containers in a cluster can be started and stopped together using the `start-containers` and `stop-containers` commands:

```
tpaexec start-containers clustername
tpaexec stop-containers clustername
```

These commands don't provision or deprovision containers, or even connect to them; they are intended to save resources when you're temporarily not using a docker cluster that you need to keep available for future use.

For a summary of the provisioned docker containers in a cluster, whether started or stopped, use the `list-containers` command:

```
tpaexec list-containers clustername
```

## 5.1 Cluster configuration

With TPA, the way to make any configuration change to a cluster is to edit `config.yml` and run the `provision/deploy/test` cycle. The process is carefully designed to be idempotent, and to make changes only in response to a change in the configuration or a change on the instances.

The *Cluster configuration* command will generate a sensible `config.yml` file for you, but it covers only the most common topology and configuration options. If you need something beyond the defaults, or you need to make changes after provisioning the cluster, you will need to edit `config.yml` anyway.

This page is an overview of the configuration mechanisms available. There's a separate page with more details about the specific *Instance configuration*.

### 5.1.1 `config.yml`

Your `config.yml` file is a **YAML format** text file that represents all aspects of your desired cluster configuration. Here's a minimal example of a cluster with two instances:

```
cluster_name: speedy

cluster_vars:
  postgres_version: 14

instances:
- node: 1
  Name: one
  role: primary
  platform: docker
  vars:
    ansible_user: root
    x: 42
```

(continues on next page)

(continued from previous page)

```

- node: 2
  Name: two
  role: replica
  platform: docker
  upstream: one
  vars:
    ansible_user: root
    x: 53

```

These three definitions are central to your cluster configuration. The file may contain many other definitions (including platform-specific details), but the list of `instances` with `vars` set either for one instance or for the whole cluster are the basic building blocks of every TPA configuration.

### All *Cluster configuration*

options translate to `config.yml` variables in some way. A single option may affect several variables (e.g., `--bdr-version` could set `postgres_version`, `edb_repositories`, `extra_postgres_extensions`, and so on), but you can always accomplish with an editor what you could by running the command.

In terms of YAML syntax, `config.yml` as a whole represents a hash with keys such as `cluster_vars` and `instances`. **You must ensure that each key is defined only once.** If you were to inadvertently repeat the `cluster_vars`, say, the second definition would completely override the former, and your next deployment could make unintended changes because of missing (shadowed) variables.

TPA checks the consistency of the overall cluster topology (for example, if you declare an instance with the role “replica”, you must also declare the name of its upstream instance, and that instance must exist), but it will not prevent you from setting any variable you like on the instances. You must exercise due caution, and try out changes in a test environment before rolling them out into production.

## 5.1.2 Variables

In Ansible terminology, most configuration settings are “inventory variables”—TPA will translate `cluster_vars` into `group_vars` (that apply to the cluster as a whole) and each instance’s `vars` into `host_vars` in the inventory during provisioning, and deployment will use the inventory values. After you change `config.yml`, **you must remember to run** `tpaexec provision` **before** `tpaexec deploy`.

Any variable can be set for the entire cluster, or an individual host, or both; host variables override group variables. In practice, setting `x: 42` in `cluster_vars` is no different from setting it in every host’s `vars`. A host that needs `x` during deployment will see the value 42 either way. A host will always see the most specific value, so it is convenient to set some default value for the group and override it for specific instances as required.

Whenever possible, defining variables in `cluster_vars` and overriding them for specific instances results in a concise configuration that is easier to review and change (less repetition). Beyond that, it’s up to you to decide whether any given setting makes more sense as a group or host variable.

### 5.1.3 Cluster variables

The keys under `cluster_vars` may map to any valid YAML type, and will be translated directly into group variables in the Ansible inventory:

```

cluster_vars:
  postgres_version: 14
  edb_repositories:

```

(continues on next page)

(continued from previous page)

```

- enterprise
- postgres_distributed
postgres_conf_settings:
  bdr.trace_replay: true

```

!!!Warning

Any variables used in templates must be defined at the top-level of `config.yml` (the same level as the `cluster_name` variable)

Please refer to the following example below.

In this case, `tpaexec` provision will write three variables (a string, a list, and a hash) to the inventory in `group_vars/tag_Cluster_name/01-cluster_name.yml`.

### 5.1.4 Instance variables

This documentation uses the term “instance variables” to refer to any variables that are defined for a specific instance in `config.yml`. For example, here’s a typical instance definition:

```

instances:
- Name: unwind
  node: 1
  backup: unkempt
  location: a
  role:
    - primary
    - bdr
  volumes:
    - device_name: root
      encrypted: true
      volume_size: 16
      volume_type: gp2
    - device_name: /dev/xvdf
      encrypted: true
      vars:
        volume_for: postgres_data
        volume_size: 64
        volume_type: gp2
  platform: aws
  type: t3.micro
  vars:
    ansible_user: ec2-user
    postgres_conf_directory: /opt/postgres/conf

```

The variables defined in this instance’s `vars` will all become host variables in the inventory, but all host vars in the inventory do not come from `vars` alone. Some other instance settings, including `platform`, `location`, `volumes`, and `role` are also copied to the inventory as host vars (but you cannot define these settings under `vars` or `cluster_vars` instead).

The settings outside `vars` may describe the properties of the instance (e.g., `Name` and `node`) or its place in the topology of the cluster (e.g., `role`, `backup`) or they may be platform-specific attributes (e.g., instance `type` and `volumes`). Other than knowing that they cannot be defined under `vars`, it is rarely necessary to distinguish between these instance “settings” and instance “variables”.

In this case, `tpaexec provision` will write a number of host variables to the inventory in `host_vars/unwind/01-instance_vars.yml`.

### 5.1.5 instance\_defaults

This is a mechanism to further reduce repetition in `config.yml`. It is most useful for instance settings that cannot be defined as `cluster_vars`. For example, you could write the following:

```
instance_defaults:
  platform: aws
  type: t3.micro
  tags:
    AWS_ENVIRONMENT_SPECIFIC_TAG_KEY: some_mandated_value

instances:
- node: 1
  Name: one

- node: 2
  Name: two

- ...
```

Whatever you specify under `instance_defaults` serves as the default for every entry in `instances`. In this example, it saves spelling out the `platform` and `type` of each instance, and makes it easier to change all your instances to a different type. If any instance specifies a different value, it will of course take precedence over the default.

It may help to think of `instance_defaults` as being a macro facility to use in defining `instances`. What is ultimately written to the inventory comes from the (expanded) definition of `instances` alone. If you're trying to decide whether to put something in `cluster_vars` or `instance_defaults`, it probably belongs in the former unless it *cannot* be defined as a variable (e.g., `platform` or `type`), which is true for many platform-specific properties (such as AWS resource tags) that are used only in provisioning, and not during deployment.

The `instance_defaults` mechanism does nothing to stop you from using it to fill in the `vars` for an instance (default hash values are merged with any hash specified in the `instances` entry). However, there is no particular advantage to doing this rather than setting the same default in `cluster_vars` and overriding it for an instance if necessary. When in doubt, use `cluster_vars`.

### 5.1.6 Locations

You can also specify a list of `locations` in `config.yml`:

```
locations:
- Name: first
  az: eu-west-1a
  region: eu-west-1
  subnet: 10.33.110.128/28

- Name: second
  az: us-east-1b
  region: us-east-1
  subnet: 10.33.75.0/24
```

(continues on next page)

(continued from previous page)

```
instances:
- node: 1
  Name: one
  location: first
...
```

If an instance specifies `location: first` (or `location: 0`), the settings under that location serve as defaults for that instance. Again, just like `instance_defaults`, an instance may override the defaults that it inherits from its location. And again, you can use this feature to fill in `vars` for an instance. This can be useful if you have some defaults that apply to only half your instances, and different values for the other half (as with the platform-specific settings in the example above).

Locations represent a collection of settings that instances can “opt in” to. You can use them to stand for different data centres, AWS regions, Docker hosts, or something else entirely. TPA does not expect or enforce any particular interpretation.

## 5.2 Instance configuration

This page presents an overview of the various controls that TPA offers to customise the deployment process on cluster instances, with links to more detailed documentation.

Before you dive into the details of deployment, it may be helpful to read [Cluster configuration](#) to understand how cluster and instance variables and the other mechanisms in `config.yml` work together to allow you to write a concise, easy-to-review configuration.

### 5.2.1 System-level configuration

The first thing TPA does is to ensure that Python is bootstrapped and ready to execute Ansible modules (a distribution-specific process). Then it completes various system-level configuration tasks before moving on to [Postgres flavour and version](#) below.

- [Distribution support](#)
- [Setting up the TPA Python environment \(preferred\\_python\\_version\)](#)
- [Environment variables \(e.g., https\\_proxy\)](#)

### Package repositories

You can use the [pre-deploy](#)

to execute tasks before any package repositories are configured.

- [Configuring YUM repositories](#)

(for RHEL, Rocky and AlmaLinux)

- [Configuring APT repositories](#)

(for Debian and Ubuntu)

- [Configuring EDB Repos 2.0 repositories](#)

(on any system)

- [Creating and using a local repository](#)

(to ship packages to target instances)

You can use the *post-repo*

to execute tasks after package repositories have been configured (e.g., to correct a problem with the repository configuration before installing any packages).

### Package installation

Once the repositories are configured, packages are installed at various stages throughout the deployment, beginning with a batch of system packages:

- *Installing packages*

(e.g., `acl`, `openssl`, `sysstat`)

Postgres and other components (e.g., Barman, `repmgr`, `pgbouncer`) will be installed separately according to the cluster configuration; these are documented in their own sections below.

### Other system-level tasks

- *Filesystem configuration* (including RAID, LUKS setup)
- *Uploading artifacts* (files, directories, tar archives)
- *Setting `sysctl` values*
- *Configuring `/etc/hosts`*
- *Managing SSH host keys*
- *Locale*

### Skipping deployment completely

To prevent TPA from doing any part of the deployment process on an instance - in other words, if you want TPA to provision the instance and then leave it alone - set the `provision_only` setting for the instance to `true` in `config.yml`. This setting will make TPA omit the instance entirely from the inventory which `tpaexec deploy` sees.

## 5.2.2 Postgres

Postgres configuration is an extended process that goes hand-in-hand with setting up other components like `repmgr` and `pgbouncer`. It begins with installing Postgres itself.

### Version selection

Use the *Software versions* to select a Postgres flavour and version, or set `postgres_version` in `config.yml` to specify which Postgres major version you want to install.

That's all you really need to do to set up a working cluster. Everything else on this page is optional. You can control every aspect of the deployment if you want to, but the defaults are carefully tuned to give you a sensible cluster as a starting point.

### Installation

The default `postgres_installation_method` is to install packages for the version of Postgres you selected, along with various extensions, according to the architecture's needs:

- *Installing Postgres-related packages*

(e.g., `pglogical`, `BDR`, etc.)

- *Postgres source installation*

(for development and testing)

Whichever installation method you choose, TPA can give you the same cluster configuration with a minimum of effort.

### Configuration

- *The postgres Unix user*
- *Running initdb*
- *pg\_hba.conf*
- *pg\_ident.conf*
- *postgresql.conf*

You can use the *postgres-config*

to execute tasks after the Postgres configuration files have been installed (e.g., to install additional configuration files).

Once the Postgres configuration is in place, TPA will go on to install and configure other components such as Barman, repmgr, pgbouncer, and haproxy, according to the details of the architecture.

### 5.2.3 Other components

- *barman-pre-config*
- *Repmgr redirect pgbouncer*
- *etcd*
- *Patroni cluster management commands*
- *Configuring haproxy*
- *Configuring HARP*
- *Configuring EFM*

### Configuring and starting services

TPA will now install systemd service unit files for each service. The service for Postgres is named `postgres.service`, and can be started or stopped with `systemctl start postgres`.

In the first deployment, the Postgres service will now be started. If you are running `tpaexec deploy` again, the service may be reloaded or restarted depending on what configuration changes you may have made. Of course, if the service is already running and there are no changes, then it's left alone.

In any case, Postgres will be running at the end of this step.

### 5.2.4 After starting Postgres

- *Creating Postgres users*
- *Creating Postgres tablespaces*
- *Creating Postgres databases*
- *pglogical*
- *Configuring .pgpass*

You can use the *postgres-config-final*

to execute tasks after the post-startup Postgres configuration has been completed (e.g., to perform SQL queries to create objects or load data).

- *EDB Postgres Distributed configuration*

You can use the *post-deploy*

to execute tasks after the deployment process has completed.

## 5.3 Building from source

!!!Warning

This option is intended for developers and advanced users. Only software built and tested by EDB is supported by EDB. Please refer to [Self-Managed Supported Open Source Software](#) .

You can even combine packaged installations of certain components with source builds of others. For example, you can install Postgres from packages and compile pglogical and PGD from source, but package dependencies would prevent installing pglogical from source and PGD from packages.

Source builds are meant for use in development, testing, and for support operations.

### 5.3.1 Quickstart

Spin up a cluster with 2ndQPostgres, pglogical3, and bdr all built from stable branches:

```
tpaexec configure ~/clusters/speedy -a BDR-Always-ON
--layout bronze
--harp-consensus-protocol etcd
--install-from-source
  2ndqpostgres:2QREL_13_STABLE_dev
  pglogical3:REL3_7_STABLE
  bdr3:REL3_7_STABLE
```

As above, but set up a cluster that builds 2ndQPostgres source code from the official git repository and uses the given local work trees to build pglogical and BDR. This feature is specific to Docker:

```
tpaexec configure ~/clusters/speedy
--architecture BDR-Always-ON --layout bronze
--harp-consensus-protocol etcd
--platform docker
--install-from-source 2ndqpostgres:2QREL_13_STABLE_dev
  pglogical3 bdr3
--local-source-directories
  pglogical3:~/src/pglogical bdr3:~/src/bdr
```

After deploying your cluster, you can use `tpaexec deploy ... --skip-tags build-clean` on subsequent runs to reuse build directories. (Otherwise the build directory is emptied before starting the build.)

Read on for a detailed explanation of how to build Postgres, pglogical, BDR, and other components with custom locations and build parameters.

### 5.3.2 Configuration

There are two aspects to configuring source builds.

If you just want a cluster running a particular combination of sources, run `tpaexec configure` to generate a configuration with sensible defaults to download, compile, and install the components you select. You can build Postgres or Postgres Extended, pglogical, and BDR, and specify branch names to build from, as shown in the examples above.

The underlying mechanism is capable of much more than the command-line options allow. By editing `config.yml`, you can clone different source repositories, change the build location, specify different configure or build parameters, redefine the build commands entirely, and so on. You can, therefore, build things other than Postgres, pglogical, and BDR.

The available options are documented here:

- [Postgres source installation](#)
- [Building and installing from source](#)

### 5.3.3 Local source directories

You can use TPA to provision Docker containers that build Postgres and/or extensions from your local source directories instead of from a Git repository.

Suppose you're using `--install-from-source` to declare what you want to build:

```
tpaexec configure ~/clusters/speedy
--architecture BDR-Always-ON --layout bronze
--harp-consensus-protocol etcd
--platform docker
--install-from-source 2ndqpostgres:2QREL_13_STABLE_dev
pglogical3:REL3_7_STABLE bdr3:REL3_7_STABLE
...
```

By default, this will clone the known repositories for Postgres Extended, pglogical3, and bdr3, check out the given branches, and build them. But you can add `--local-source-directories` to specify that you want the sources to be taken directly from your host machine instead:

```
tpaexec configure ~/clusters/speedy
--architecture BDR-Always-ON --layout bronze
--harp-consensus-protocol etcd
--platform docker
--install-from-source 2ndqpostgres:2QREL_13_STABLE_dev
pglogical3 bdr3
--local-source-directories
pglogical3:~/src/pglogical bdr3:~/src/bdr
...
```

This configuration will still install Postgres Extended from the repository, but it obtains pglogical3 and bdr3 sources from the given directories on the host. These directories are bind-mounted read-only into the Docker containers at the same locations where the git repository would have been cloned to, and the default (out-of-tree) build proceeds as usual.

If you specify a local source directory for a component, you cannot specify a branch to build (cf. `pglogical3:REL3_7_STABLE` vs. `pglogical3` for `--install-from-source` in the examples above). The source directory is mounted read-only in the containers, so TPA cannot do anything to change it—neither `git pull`, nor `git checkout`. You get whichever branch you have checked out locally, uncommitted changes and all.

Using `--local-source-directories` includes a list of Docker volume definitions in `config.yml`:

**local\_source\_directories:**

- /home/ams/src/pglogical:/opt/postgres/src/pglogical:ro
- /home/ams/src/bdr:/opt/postgres/src/bdr:ro
- ccache-bdr\_src\_36-20200828200021:/root/.ccache:rw

**ccache**

TPA installs ccache by default for source builds of all kinds. When you are using a Docker cluster with local source directories, by default a new Docker volume is attached to the cluster's containers to serve as a shared ccache directory. This volume is completely isolated from the host, and is removed when the cluster is deprovisioned.

Use the `--shared-ccache /path/to/host/ccache` configure option to specify a longer-lived shared ccache directory. This directory will be bind-mounted r/w into the containers, and its contents will be shared between the host and the containers.

(By design, there is no way to install binaries compiled on the host directly into the containers.)

### 5.3.4 Rebuilding

After deploying a cluster with components built from source, you can rebuild those components quickly without having to rerun `tpaexec deploy` by using the `tpaexec rebuild-sources` command. This will run `git pull` for any components built from git repositories on the containers, and rebuild all components.

## 5.4 TPA hooks

TPA can set up fully-functional clusters with no user intervention, and already provides a broad variety of *Instance configuration*, including custom repositories and packages, custom Postgres configuration (both `pg_hba.conf` and `postgresql.conf`), and so on.

You can write hook scripts to address specific needs that are not met by the available configuration settings. Hooks allow you to execute arbitrary Ansible tasks during the deployment.

Hooks are the ultimate extension mechanism for TPA, and there is no limit to what you can do with them. Please use them with caution, and keep in mind the additional maintenance burden you are taking on. The TPA developers have no insight into your hook code, and cannot guarantee compatibility between releases beyond invoking hooks at the expected stage.

### 5.4.1 Summary

If you create files with specific names under the `hooks` subdirectory of your cluster directory, TPA will invoke them at various stages of the deployment process, as described below.

```
mkdir ~/clusters/speedy/hooks
cat > ~/clusters/speedy/hooks/pre-deploy.yml
```

```
- --
- debug: msg="hello world!"
```

Hook scripts are invoked with `include_tasks`, so they are expected to be YAML files containing a list of Ansible tasks (not a playbook, which contains a list of plays). Unless otherwise documented below, hooks are unconditionally executed for all hosts in the deployment.

## 5.4.2 General-purpose hooks

### pre-deploy

TPA invokes `hooks/pre-deploy.yml` immediately after bootstrapping Python—but before doing anything else like configuring repositories and installing packages. This is the earliest stage at which you can execute your own code.

You can use this hook to set up custom repository configuration, beyond what you can do with *Configuring APT repositories* or *Configuring YUM repositories*.

### post-repo

TPA invokes `hooks/post-repo.yml` after configuring package repositories. You can use it to make corrections to the repository configuration before beginning to install packages.

### pre-initdb

TPA invokes `hooks/pre-initdb.yml` before deciding whether or not to *Running initdb* if it does not exist. You should not ordinarily need to use this hook (but if you use it to create `PGDATA` yourself, then TPA will skip `initdb`).

### postgres-config

TPA invokes `hooks/postgres-config.yml` after generating Postgres configuration files, including `pg_hba.conf` and the files in `conf.d`, but before the server has been started.

You can use this hook, for example, to create additional configuration files under `conf.d`.

### postgres-config-final

TPA invokes `hooks/postgres-config-final.yml` after starting Postgres and creating users, databases, and extensions. You can use this hook to execute SQL commands, for example, to perform custom extension configuration or create database objects.

### barman-pre-config

TPA invokes `hooks/barman-pre-config.yml` after installing Barman and setting up Barman users, but before generating any Barman configuration.

You can use this hook, for example, to perform any tasks related with Barman certificate files or mount points.

### efm-pre-config

TPA invokes `hooks/efm-pre-config.yml` after installing efm, creating its configuration directory, and setting up the efm user, but before generating any efm configuration.

An example use of this hook is to install efm helper scripts.

### efm-post-config

TPA invokes `hooks/efm-post-config.yml` after finishing the efm configuration step.

An example use of this hook is for additional steps that need to be taken into account before exiting the efm configuration process.

### **harp-config**

TPA invokes `hooks/harp-config.yml` after generating HARP configuration files, but before the HARP service has been started.

You can use this hook, for example, to perform any customizations to the HARP proxy that are not provided by the built-in interface of TPA.

Please note that this hook will be run in any node that installs HARP packages, including PGD nodes.

### **pgd-proxy-config**

TPA invokes `hooks/pgd-proxy-config.yml` after generating PGD PROXY configuration files, but before the PGD PROXY service has been started.

You can use this hook, for example, to perform any customizations to the PGD Proxy that are not provided by the built-in interface of TPA.

Note that this hook will run in any node that installs PGD PROXY packages.

### **post-deploy**

TPA invokes `hooks/post-deploy.yml` at the end of the deployment.

You can go on to do whatever you want after this stage.

If you use this hook to make changes to any configuration files that were generated or altered during the TPA deployment, you run the risk that the next `tpaexec deploy` will overwrite your changes (since TPA doesn't know what your hook might have done).

## **5.4.3 PGD hooks**

These hooks are specific to PGD deployments.

### **bdr-pre-node-creation**

TPA invokes `hooks/bdr-pre-node-creation.yml` on all instances before creating a PGD node on any instance for the first time. The hook will not be invoked if all required PGD nodes already exist.

### **bdr-post-group-creation**

TPA invokes `hooks/bdr-post-group-creation.yml` on all instances after creating any PGD node group on the `first_bdr_primary` instance. The hook will not be invoked if the required PGD groups already exist.

### **bdr-pre-group-join**

TPA invokes `hooks/bdr-pre-group-join.yml` on all instances after creating, changing or removing the replication sets and configuring the required subscriptions, before the node join.

You can use this hook to execute SQL commands and perform other adjustments to the replication set configuration and subscriptions that might be required before the node join starts.

For example, you can adjust the PGD witness replication set to automatically add new tables and create DDL filters in general.

## 5.4.4 Other hooks

### postgres-pre-update, postgres-post-update

The *Upgrading your cluster* command invokes `hooks/postgres-pre-update.yml` on a particular instance before it installs any packages, and invokes `hooks/postgres-post-update.yml` after the package installation is complete. Both hooks are invoked only on the instance being updated.

You can use these hooks to customise the update process for your environment (e.g., to install other packages and stop and restart services that TPA does not manage).

## 5.4.5 New hooks

EDB adds new hooks to TPA as the need arises. If your use case is not covered by the existing hooks, please contact us to discuss the matter.



## 6.1 Upgrading your cluster

The `tpaexec upgrade` command is used to upgrade the software running on your TPA cluster (`tpaexec deploy` will not perform upgrades).

(This command replaces the earlier `tpaexec update-postgres` command.)

!!!Note

TPA does not yet support using the `tpaexec upgrade` command for clusters that have shared Barman and/or shared PEM configurations, and will add this functionality in a future release.

If you make any changes to `config.yml`, the way to apply those changes is to run `tpaexec provision` followed by `tpaexec deploy`.

The exception to this rule is that `tpaexec deploy` will refuse to install a different version of a package that is already installed. Instead, you must use `tpaexec upgrade` to perform software upgrades.

The following components are able to be upgraded for **any** architecture:

- Postgres
- *Repmgr redirect pgbouncer*
- *barman-pre-config*
- *PG Backup API*
- *Configuring Postgres Enterprise Manager (PEM)*

(both server and agent)

The following components are able to be upgraded on **M1** architectures and depend on the failover manager used:

- *efm-pre-config*
- *Patroni cluster management commands*
- *etcd* (for Patroni)

- *Repmgr redirect pgbouncer*

The following components are able to be upgraded on **BDR-Always-ON/PGD-Always-ON** architectures, depending on the BDR version used:

- *pgdcli* (v1 for BDR 4, v5 for PGD5)
- *pgd-proxy-config* (only for PGD5)

!!!Note Minor version upgrades only

**``tpaexec upgrade`` does NOT support MAJOR version upgrades of Postgres and most cluster components** What TPA can upgrade is dependent on architecture:

- The M1 architecture and all applicable failover managers for M1, upgrade can perform minor version upgrades of Postgres, the corresponding failover manager (EFM, Patroni or repmgr) and any non-architecture specific components that are selected.
- With PGD architectures, upgrade will perform minor version upgrades of Postgres and the BDR extension as well as *pgd-cli* and *pgd-proxy* if they are explicitly opted-in.
- With PGD architectures, and only in combination with the *reconfigure* command, upgrade can perform major-version upgrades of the BDR extension.

Support for upgrading other cluster components was added in TPA v23.41.0 *Certain components (such as ``EFM``) are an exception. Consult the ``upgrade`` section of each component's documentation for further information.*

When upgrading, you should always use barman to take a backup before beginning the upgrade and disable any scheduled backups which would take place during the time set aside for the upgrade.

In general, TPA will proceed instance-by-instance, stopping any affected services, installing new packages, updating the configuration if needed, restarting services, and performing any runtime configuration changes, before moving on to do the same thing on the next instance. At any time during the process, only one of the cluster's nodes will be unavailable.

When upgrading a cluster to PGD-Always-ON or upgrading an existing PGD-Always-ON cluster, you can enable monitoring of the status of your proxy nodes during the upgrade by adding the option `-e enable_proxy_monitoring=true` to your `tpaexec upgrade` command line. If enabled, this will create an extra table in the bdr database and write monitoring data to it while the upgrade takes place. The performance impact of enabling monitoring is very small and it is recommended that it is enabled.

### 6.1.1 Component selection

!!!Note

Upgrading components is strictly opt-in

```
tpaexec upgrade ~/clusters/speedy
```

To select specific components to update, the `--components` flag takes a comma-separated list

```
tpaexec upgrade ~/clusters/speedy \  
--components=postgres,pgd-proxy,pgdcli,pgbouncer,pg-backup-api,barman
```

If all applicable components in the cluster should be updated, `all` can be passed to the flag

```
tpaexec upgrade --components=all
```

| component         | value                | architecture                            |
|-------------------|----------------------|-----------------------------------------|
| Barman            | barman               | all                                     |
| PEM               | pem-server,pem-agent | all                                     |
| PgBackupAPI       | pg-backup-api        | all                                     |
| PgBouncer         | pgbouncer            | all                                     |
| EFM               | efm                  | M1 with <i>failover_manager=efm</i>     |
| etcd              | etcd                 | M1 with <i>failover_manager=patroni</i> |
| Patroni           | patroni              | M1 with <i>failover_manager=patroni</i> |
| RepMgr            | repmgr               | M1 with <i>failover_manager=repmgr</i>  |
| PGD Cli           | pgdcli               | BDR-Always-ON,<br>PGD-Always-ON         |
| PGD-Proxy         | pgd-proxy            | PGD-Always-ON                           |
| Postgres,EPAS,PGE | postgres             | All                                     |
| All               | all                  | All                                     |

## 6.1.2 Package version selection

By default, `tpaexec upgrade` will update to the latest available versions of the installed packages if you did not explicitly specify any package versions (e.g., Postgres, PGD, or pglogical) when you created the cluster.

!!!Note Minor upgrade is not strictly enforced

If a desired package version is NOT provided when upgrading, TPA will install the latest available package. The minor version restriction is NOT strictly enforced during `tpaexec upgrade`. This can result in unwillingly attempting an unsupported major upgrade of a component. Thus, it is recommended to explicitly select versions for upgrade to ensure compatibility in the existing cluster. Postgres does not pose this issue since major versions are different packages altogether which stops this from happening.

In this case, you must edit `config.yml`, update the version settings, and re-run `tpaexec provision`. The update will then install the selected version of the packages. You can also update to a specific version by specifying versions on the command line as shown below:

```
tpaexec upgrade ~/clusters/speedy -vv \
--components=postgres,pgbouncer \
-e postgres_package_version="16.10*" \
-e pgbouncer_package_version="1.24*" \
-e bdr_package_version="5.9.0*" \
```

Please note that version syntax here depends on your OS distribution and package manager. In particular, yum accepts `*xyz*` wildcards, while apt only understands `xyz*` (as in the example above).

!!!Note : see limitations of using wildcards in `package_version` in

*Known issue with wildcard use .*

For PGD clusters, it is a good idea to explicitly specify exact versions for all three components (Postgres, PGD, pglogical) rather than rely on the package manager's dependency resolution to select the correct dependencies.

We strongly recommend testing the upgrade in a QA environment before running it in production.

### 6.1.3 Configuration

In certain cases, minor-version upgrades do not need changes to `config.yml`. If no `postgres_package_version` is defined in `config.yml`, when `tpaexec upgrade` is run, it will upgrade Postgres to the latest available minor-version in a graceful way (what exactly that means depends on the details of the cluster).

For control over minor-version upgrades of other components, it is recommended to ensure a specific `xxx_package_version` is specified in `config.yml` before running `tpaexec upgrade` and explicitly opting-in to upgrade specific components using the `--components=x,y,z` flag (or `--components=all` to upgrade all, as applicable to the cluster). Running `tpaexec upgrade` and opting in to upgrade some or all components WITHOUT pinning their `xxx_package_version` in `config.yml` could result in a major version upgrade of installed component packages, which TPA does not support as it may break compatibility.

Sometimes an upgrade involves additional steps beyond installing new packages and restarting services. For example, in order to upgrade from BDR4 to PGD5, one must set up new package repositories and make certain changes to the BDR node and group configuration during the process.

In such cases, where there are complex steps required as part of the process of effecting a software upgrade, `tpaexec upgrade` will perform those steps. For example, in the above scenario, it will configure the new PGD5 package repositories (which deploy would also normally do).

However, it will make only those changes that are directly required by the upgrade process itself. For example, if you edit `config.yml` to add a new Postgres user or database, those changes will not be done during the upgrade. To avoid confusion, we recommend that you `tpaexec deploy` any unrelated pending changes before you begin the software upgrade process.

### 6.1.4 Upgrading from BDR-Always-ON to PGD-Always-ON

To upgrade from BDR-Always-ON to PGD-Always-ON (that is, from BDR3/4 to PGD5), first run `tpaexec reconfigure`:

```
tpaexec reconfigure ~/clusters/speedy\  
  --architecture PGD-Always-ON\  
  --pgd-proxy-routing local
```

This command will read `config.yml`, work out the changes necessary to upgrade the cluster, and write a new `config.yml`. For details of its invocation, see [tpaexec reconfigure](#). After reviewing the changes, run `tpaexec upgrade` to perform the upgrade:

```
tpaexec upgrade ~/clusters/speedy\  

```

Or to run the upgrade with proxy monitoring enabled,

```
tpaexec upgrade ~/clusters/speedy\  
  -e enable_proxy_monitoring=true
```

`tpaexec upgrade` will automatically run `tpaexec provision`, to update the ansible inventory. The upgrade process does the following:

1. Checks that all preconditions for upgrading the cluster are met.
2. For each instance in the cluster, checks that it has the correct repositories configured and that the required postgres packages are available in them.
3. For each BDR node in the cluster, one at a time:
  - Fences the node off to ensure that harp-proxy doesn't send any connections to it.
  - Stops, updates, and restarts postgres, including replacing BDR4 with PGD5.

- Unfences the node so it can receive connections again.
  - Updates pgbouncer and pgd-cli, as applicable for this node.
4. For each instance in the cluster, updates its BDR configuration specifically for BDR v5
  5. For each proxy node in the cluster, one at a time:
    - Sets up pgd-proxy.
    - Stops harp-proxy.
    - Starts pgd-proxy.
  6. Removes harp-proxy and its support files.

### 6.1.5 Upgrading from PGD-Always-ON to PGD-X

Upgrading a PGD-Always-ON cluster to PGD-X is a **significant architectural evolution**, involving changes beyond a simple **software update**. It is a *carefully orchestrated, multi-stage process* that requires reconfiguring your cluster in distinct phases before the final software upgrade can take place. The procedure first modernizes your PGD 5 cluster's connection handling by replacing pgd-proxy with the built-in Connection Manager—a step that currently requires manual operations on the live cluster but is planned for automation in a future TPA release—and then transitions the cluster to the new PGD-X architecture.

The upgrade process transitions the cluster through three distinct states:

1. **Start:** PGD 5.9+ (PGD-Always-ON ) using PGD-Proxy
2. **Intermediate:** PGD 5.9+ (PGD-Always-ON ) now using the built-in Connection Manager
3. **Final:** PGD 6 (PGD-X Architecture )

#### Prerequisites

Before you begin, ensure you have met the following requirements:

- **Cluster Version:** Your cluster must be running PGD version 5.9 or later. If you are on an earlier 5.x version, use `tpaexec upgrade` to upgrade to the latest minor version first. See the section (#pgd-always-on) for details on minor version upgrade of a PGD-Always-ON cluster.
- **Backup:** You have a current, tested backup of your cluster.
- **Review Overrides:** You have reviewed your `config.yml` for any instance-level proxy overrides (e.g., `pgd_proxy_options`). These cannot be migrated automatically and will require manual intervention.
- **Co-hosted Proxies:** Your PGD 5 cluster must be configured with co-hosted proxies (where the `pgd-proxy` role is on the same instance as the `bdr` role). The presence of standalone proxy instances will cause the `switch2cm` command to abort. You must remove standalone proxy instances from your cluster before proceeding with the migration.

#### Stage 1: Migrating to the Built-in Connection Manager

The first stage is to reconfigure your PGD 5.9+ cluster to switch from using the external `pgd-proxy` to the modern, built-in Connection Manager. TPA provides the `tpaexec switch2cm` command to automate this migration with minimal downtime.

!!!Note Transitional State Only

This process creates a transitional PGD 5.9+ cluster state that is intended only as an intermediate step before upgrading to PGD 6. TPA does not currently support staying in PGD5.9+ with Connection Manager enabled or moving to a newer minor version of PGD5.9+ with this configuration. A future TPA release will fully support lifecycle management of PGD 5 with Connection Manager.

Run the following command to update your `config.yml` file. This adds the settings required to enable the built-in Connection Manager.

**This action only modifies the configuration file; it does not change the running state of your database cluster yet.**

Before writing the new version, `reconfigure` automatically saves a backup of the current file (e.g., `config.yml.~1~`), providing a safe restore point.

For details of its invocation, see [\*tpaexec reconfigure\*](#).

```
tpaexec reconfigure ~/clusters/speedy --enable-connection-manager
```

### Step 1.2: Switch to Connection Manager

Run the `tpaexec switch2cm` command to perform the migration from `pgd-proxy` to the built-in Connection Manager. This command automatically runs `tpaexec provision` to update the Ansible inventory, then switches all nodes with minimal downtime:

```
tpaexec switch2cm ~/clusters/speedy
```

The `switch2cm` command performs the following operations:

1. Updates the Ansible inventory with Connection Manager settings
2. For each node:
  - Fences the node to prevent new connections
  - Restarts PostgreSQL to load the Connection Manager configuration
  - Stops the `pgd-proxy` service
  - Restarts PostgreSQL again to allow Connection Manager to bind ports
  - Waits for Connection Manager to start listening
  - Unfences the node and verifies connectivity

This process follows the official [EDB Connection Manager Migration procedure](#).

### Stage 1 Complete

At the end of this stage, you will have a PGD cluster running with the built-in Connection Manager. This is an intermediate state, and you should proceed directly to Stage 2. While `tpaexec upgrade` for minor version upgrades is **not supported** in this intermediate state, we also advise against running `tpaexec deploy` until the upgrade to PGD 6 is complete.

### Stage 2: Upgrading the Architecture to PGD-X

Once your cluster is running with the Connection Manager, you can proceed with the final configuration step to prepare for the PGD 6 upgrade.

!!!Note

You **must** start this process from a cluster that has successfully completed Stage 1 and is running with the built-in Connection Manager.

Run the following command to update your `config.yml` for the new architecture. This changes the cluster architecture type, sets the BDR version to 6, and removes any obsolete legacy settings.

**This action only modifies the configuration file; it does not change the running state of your database cluster yet.**

```
tpaexec reconfigure ~/clusters/speedy --architecture PGD-X
```

## Step 2.2: Perform the Software Upgrade

After reviewing the final changes in `config.yml`, you can now run the standard `tpaexec upgrade` command. This will perform the software upgrade on all nodes, bringing your cluster to PGD 6.

```
tpaexec upgrade ~/clusters/speedy
```

Or to run the upgrade with proxy monitoring enabled,

```
tpaexec upgrade ~/clusters/speedy\  
-e enable_proxy_monitoring=true
```

`tpaexec upgrade` will automatically run `tpaexec provision`, to update the ansible inventory. The upgrade process does the following:

1. Checks that all preconditions for upgrading the cluster are met.
2. For each instance in the cluster, checks that it has the correct repositories configured and that the required postgres packages are available in them.
3. For each BDR node in the cluster, one at a time:
  - Fences the node off so there are no connections to it.
  - Stops, updates, and restarts postgres, including replacing PGD5 with PGD6.
  - Unfences the node so it can receive connections again.
  - Updates pgbouncer and pgd-cli, as applicable for this node.
4. Applies BDR configuration specifically for BDR v6

## Upgrade Complete

Your cluster is now running PGD 6 with the PGD-X architecture and is fully manageable with both `tpaexec deploy` and `tpaexec upgrade` as usual.

### 6.1.6 PGD-S or PGD-X

When upgrading an existing PGD6 (PGD-S or PGD-X) cluster to the latest available software versions, the upgrade process does the following:

1. Checks that the cluster is healthy and that the nodes are listening on the configured ports.
2. Checks that the nodes to be upgraded have their repositories configured and updated, including local repositories.
3. Checks that updated packages can be installed
4. Upgrade each BDR node in the cluster one at a time:

**Important:** To ensure high availability, if the write leader is among the nodes being upgraded, it will be the very last node to be upgraded.

- Fences the node off so it doesn't accept connections
- Stops postgres
- Updates postgres and PGD packages
- Unfences the node so it can receive connections again

- Checks that the BDR cluster has re-established Raft consensus
  - Checks that the upgraded node is listening on the configured ports
5. Re-runs the cluster health checks
  6. Outputs information about the upgraded packages

### **6.1.7 PGD-Always-ON**

When upgrading an existing PGD-Always-ON (PGD5) cluster to the latest available software versions, the upgrade process does the following:

1. Checks that all preconditions for upgrading the cluster are met, including that it is not a shared PEM or shared Barman cluster.
2. Runs pre-upgrade health checks for all components, as applicable to the cluster, including that no Barman backup is underway (this stops the WAL-receiver)
3. For each instance in the cluster, checks that it has the correct repositories configured and that the required postgres packages are available in them.
4. Checks that all selected components are able to be updated to the desired version (if a package version is provided)
5. For each BDR node in the cluster, one at a time:
  - Fences the node off to ensure that pgd-proxy doesn't send any connections to it.
  - Stops, updates, and restarts postgres.
  - Unfences the node so it can receive connections again.
  - Updates pgd-proxy and pgd-cli software (if explicitly opted-in)
6. For the applicable nodes in the cluster, updates pgbouncer, barman, pg-backup-api, and PEM agents/PEM server (according to the node's roles)
7. Starts the Barman WAL-receiver if required and runs post-upgrade health checks for all components (as applicable to the cluster)

### **6.1.8 BDR-Always-ON**

For BDR-Always-ON clusters, the upgrade process goes through the cluster instances one by one and does the following:

1. Checks that all preconditions for upgrading the cluster are met, including that it is not a shared PEM or shared Barman cluster.
2. Runs pre-upgrade health checks for all components, as applicable to the cluster, including that no Barman backup is underway (this stops the WAL-receiver)
3. For each instance in the cluster, checks that it has the correct repositories configured and that the required postgres packages are available in them.
4. Tell haproxy the server is under maintenance.
5. If the instance was the active server, request pgbouncer to reconnect and wait for active sessions to be closed.
6. Stop Postgres, update Postgres, etcd, and pgdcli (if applicable and opted-in) packages and restart Postgres.
7. Finally, mark the server as "ready" again to receive requests through haproxy.
8. For the applicable nodes in the cluster, updates pgbouncer, barman, pg-backup-api, and PEM agents/PEM server (according to the node's roles)

9. Starts the Barman WAL-receiver if required and runs post-upgrade health checks for all components (as applicable to the cluster)

PGD logical standby or physical replica instances are updated without any haproxy or pgbouncer interaction. Non-Postgres instances in the cluster are left alone.

### 6.1.9 M1

For M1 clusters, `upgrade` will first update the streaming replicas and witness nodes when applicable, then perform a *tpaexec switchover*

from the primary to one of the upgraded replicas, update the primary, and switchover back to the initial primary node.

1. Checks that all preconditions for upgrading the cluster are met, including that it is not a shared PEM or shared Barman cluster.
2. Runs pre-upgrade health checks for all components, as applicable to the cluster, including that no Barman backup is underway (this stops the WAL-receiver)
3. For each instance in the cluster, checks that it has the correct repositories configured and that the required postgres packages are available in them.
4. Update Postgres on the streaming replicas and witness. nodes (when applicable)
5. Perform a *tpaexec switchover* from the primary to one of the upgraded replicas
6. Update Postgres on the primary
7. Switchover back to the initial primary node.
8. For the applicable nodes in the cluster, updates pgbouncer, barman, pg-backup-api, and PEM agents/PEM server (according to the node's roles)
9. Starts the Barman WAL-receiver if required and runs post-upgrade health checks for all components (as applicable to the cluster)

### 6.1.10 Controlling the upgrade process

You can control the order in which the cluster's instances are upgraded by defining the `update_hosts` variable:

```
tpaexec upgrade ~/clusters/speedy \
-e update_hosts=quirk,keeper,quaver
```

This may be useful to minimise lead/shadow switchovers during the upgrade by listing the active PGD primary instances last, so that the shadow servers are upgraded first.

If your environment requires additional actions, the *TPA hooks*

allow you to execute custom Ansible tasks before and after the package installation step.

### 6.1.11 Upgrading a Subset of Nodes

You can perform a rolling upgrade on a subset of instances by setting the `update_hosts` variable. However, support for this feature varies by architecture.

- For the **M1** architecture, this feature is fully supported for `repmgr` and `Patroni` managed clusters. `EFM` managed clusters respect the `update_hosts` list for all components *except* `EFM`. All data nodes will upgrade their `EFM` version regardless of the nodes specified in `update_hosts` as `EFM` does not support clusters running different versions across data nodes.
- For **PGD-Always-ON/BDR-Always-ON**, this is supported **only** during minor version upgrades.

## Best Practice for PGD-Always-ON/BDR-Always-ON

When performing a minor upgrade on a subset of PGD nodes, it is highly recommended to update the **RAFT leader nodes last**. This strategy avoids potential issues with post-upgrade checks while the cluster is running mixed versions of BDR.

## 6.2 tpaexec switchover

The `tpaexec switchover` command performs a controlled switchover between a primary and a replica in a *MI*. After you run this command, the selected replica is promoted to be the new primary, the former primary becomes a new replica, and any other replicas in the cluster will be reconfigured to follow the new primary.

The command checks that the cluster is healthy before switching roles, and is designed to be run without having to shut down any `repmgr` services beforehand.

(This is equivalent to running `repmgr standby switchover` with the `--siblings-follow` option.)

### 6.2.1 Example

This command will make `replicaname` be the new primary in `~/clusters/speedy`:

```
tpaexec switchover ~/clusters/speedy replicaname
```

### 6.2.2 Architecture options

This command is applicable only to *MI*

that have a single writable primary instance and one or more read-only replicas.

For BDR-Always-ON clusters, use the *BDR/HAProxy server pool management* to perform maintenance on PGD instances.

### 6.2.3 Repmgr redirect pgbouncer

When using `repmgr` as failover manager, `pgbouncer` as connection pooler and setting `repmgr_redirect_pgbouncer: true`, `switchover` command ensures that the `pgbouncer` instances are redirected to the new primary node.

```
# switchover to a replica
tpaexec switchover <cluster_name> <replica_name>
# revert to initial primary
tpaexec switchover <cluster_name> <initial_primary_name> -e"revert_redirect=true"
```

## 6.3 BDR/HAProxy server pool management

The `tpaexec pool-disable-server` and `pool-enable-server` commands allow a PGD instance in a *BDR-Always-ON* to be temporarily removed from the HAProxy active server pool for maintenance, and then added back afterwards.

These commands follow the same process as *Upgrading your cluster* by default, so `pool-disable-server` will wait for active transactions against a PGD instance to complete and for `pgbouncer` to direct new connections to another instance before completing. Use the `--nowait` option if you don't want to wait for active sessions to end.

Running `pool-disable-server` immediately followed by `pool-enable-server` on an instance will have the effect of moving all active traffic to a different instance (in essence, a switchover). This allows you to run online maintenance tasks such as long-running `VACUUM` commands, while maintaining instance availability.

If there are multiple HAProxy servers configured with the same set of `haproxy_backend_servers`, this command will remove or add the given server to the pool of every relevant proxy in parallel.

### 6.3.1 Example

```
tpaexec pool-disable-server ~/clusters/clockwork orange # --nowait

# HAProxy will no longer direct any traffic to the PGD instance named
# orange, so you can perform maintenance on it (e.g., run `tpaexec
# rehydrate`).

tpaexec pool-enable-server ~/clusters/clockwork orange
```

When you remove an instance from the server pool, HAProxy will not direct any traffic to it, even if the other instance(s) in the pool fail. You must remember to add the server back to the active pool once the maintenance activity is concluded.

## 6.4 tpaexec rehydrate

The `tpaexec rehydrate` command rebuilds AWS EC2 instances with an updated machine image (AMI), and allows for the rapid deployment of security patches and OS upgrades to a cluster managed by TPA.

Given a new AMI with all the required changes, this command terminates an instance, replaces it with a newly-provisioned instance that uses the new image, and attaches the data volumes from the old instance before recreating the configuration of the server exactly (based on `config.yml`).

Publishing up-to-date images and requiring servers to be rebuilt from scratch on a regular schedule is an alternative to allowing a fleet of servers to download and install individual security updates themselves. It makes it simpler to track the state of each server at a glance, and discourages any manual changes to individual servers (they would be wiped out during the instance replacement).

TPA makes it simple to minimise disruption to the cluster as a whole during the rehydration, even though the process must necessarily involve downtime for individual servers as they are terminated and replaced. On a *MI*, you can rehydrate the replicas first, then use *tpaexec switchover*

to convert the primary to a replica before rehydrating it. On *BDR-Always-ON*, you can *BDR/HAProxy server pool management* before rehydrating it, then add it back afterwards.

If you just want to install minor-version updates to Postgres and associated components, you can use the *Upgrading your cluster* command instead.

### 6.4.1 Prerequisites

To be able to rehydrate an instance, you must specify `delete_on_termination: no` and `attach_existing: yes` for each of its data volumes in `config.yml`. (The new instance will necessarily have a new EBS root volume.)

By default, when you terminate an EC2 instance, the EBS volumes attached to it are also terminated. In this case, since we want to reattach them to a new instance, we must disable `delete_on_termination`. Setting `attach_existing` makes TPA search for old volumes when provisioning a new instance and, if found, attach them to the instance after it's running.

Do not stop or terminate the old instance manually; the `tpaexec rehydrate` command will do this after verifying that the instance can be safely rehydrated.

## 6.4.2 Example

Let's assume you have an AWS cluster configuration in `~/clusters/night`.

### Change the configuration

First, you must edit `config.yml` and specify the new AMI. For example:

```
ec2_ami:
  Name: RHEL-8.3_HVM-20210209-x86_64-0-Hourly2-GP2
  Owner: 309956199498
```

Check that `delete_on_termination` is disabled for each data volume. If the parameter is not present, you can check its value through the AWS EC2 management console. Click on 'Instances', select an instance, then open the 'Description' tab and scroll down to 'Block devices', and click on an EBS volume. If the "Delete on termination" flag is set to true, you can [Appendix](#). Also check `attach_existing` and set it to `yes` if it isn't set already.

Here's an example with both attributes correctly set:

```
instances:
- node: 1
  Name: vlad
  subnet: 10.33.14.0/24
  role: primary
  volumes:
  - device_name: /dev/xvdf
    volume_type: gp2
    volume_size: 16
    attach_existing: yes
    delete_on_termination: false
  vars:
    volume_for: postgres_data
    mountpoint: /var/lib/pgsql
```

(Note that volume parameters may be set in `instance_defaults` as well as under specific instances. Search for `volumes:` and make sure all of the relevant volumes have these two attributes set.)

### Start the rehydration

Here's the syntax for the rehydrate command:

```
tpaexec rehydrate ~/clusters/night instancename
```

You can specify a single instance name or a comma-separated list of instance names (but you cannot rehydrate all of the instances in the cluster at once).

The command will first check that every non-root EBS volume attached to the instance (or instances) being rehydrated has the `delete_on_termination` flag set to false. If this is not the case, it will stop with an error before any instance is terminated.

If the volume attributes are set correctly, the command will first terminate each of the instances, then run provision and deploy to replace them with new instances using the new AMI.

### 6.4.3 Rehydrate in phases

In order to maintain cluster continuity, we recommend rehydrating the cluster in phases.

For example, in a *MI* with a primary instance, two replicas, and a Barman backup server, you could rehydrate the Barman instance and one replica first, then another replica, then *tpaexec switchover* from the primary to one of the rehydrated replicas, rehydrate the former primary, and (optionally), switchover back to the original primary. This sequence ensures that one primary and one replica are always available.

### 6.4.4 Appendix

#### Using awscli to change volume attributes

First, find the instance and EBS volume in the AWS management console. Click on ‘Instances’, select an instance, open the ‘Description’ tab and scroll down to ‘Block devices’, and select an EBS volume. To disable `delete_on_termination`, run the following command after substituting the correct values for the `--region`, `--instance-id`, and block device name:

```
aws ec2 modify-instance-attribute \
  --region eu-west-1 --instance-id i-XXXXXXXXXXXXXXXXX \
  --block-device-mappings \
    [{"DeviceName": "/dev/xvdf", "Ebs": {"DeleteOnTermination": false}}]
```

Do this for each of the data volumes for the instance, and after a brief delay, you should be able to see the changes in the management console, and `tpaexec rehydrate` will also detect that the instance can be safely rehydrated.



### 7.1 TPA and Ansible Tower/Ansible Automation Platform

TPA supports deployments via RedHat Ansible Automation Platform (AAP). The support, as detailed below, works by allowing you to run `deploy` and `upgrade` steps on AAP. Before you can run `deploy` or `upgrade` (later), you will need to run configuration (`configure` command) and provisioning (`provision` command) on a separate standalone machine that has tpa packages installed. Once you have run `configure` and `provision` on this standalone machine with suitable options, you can then import the resulting cluster directory on AAP. Support is limited to bare-metal platforms.

#### 7.1.1 AAP initial setup

Before TPA can use AAP to deploy clusters, you need to perform this initial setup.

##### Add TPA Execution Environment image (admin)

Starting with version 2.4, AAP uses container images to run Ansible playbooks. These containers, called Execution Environments (EE), bundle dependencies required by playbooks to run successfully. As a consequence, this means that in order to, resolve and use all required TPA dependencies, you will need an EE that includes TPA so your AAP can use it when running deployments and upgrades.

EDB customers can reach out to EDB Support for help with EE.

As an AAP admin, create an entry in your available EE list that points to your TPA enabled EE image.

##### Create the EDB\_SUBSCRIPTION\_TOKEN credential type (admin)

As an AAP admin, create the custom credential type `EDB_SUBSCRIPTION_TOKEN` to hold your EDB subscription access token:

1. Go to the Credentials Type page in the AAP UI.
2. Set the **Name** field to `EDB_SUBSCRIPTION_TOKEN`.
3. Paste the following into the **Input Configuration** field:

```
fields:
- id: tpa_edb_sub_token
  type: string
  label: EDB_SUBSCRIPTION_TOKEN
  secret: true
required:
- tpa_edb_sub_token
```

4. Paste the following into the **Injector Configuration** field:

```
env:
  EDB_SUBSCRIPTION_TOKEN: {{ tpa_edb_sub_token }}
```

5. Save the changes.
6. Create a credential using the newly added type `EDB_SUBSCRIPTION_TOKEN`.

## 7.1.2 Setting up a cluster

Perform the initial steps on a workstation with the `tpaexec` packages installed.

### On the TPA workstation

#### Configure

Run the `tpaexec configure` command, including these options: `--platform bare`, `--use-ansible-tower`, `--tower-git-repository`

```
tpaexec configure <clustername> \
  --platform bare \
  --use-ansible-tower https://aac.example.com \
  --tower-git-repository ssh://git@git.example.com/example \
  --hostnames-from <hostnamefile> \
  --architecture PGD-Always-ON \
  --pgd-proxy-routing local \
  --postgresql 16
```

`--use-ansible-tower` expects the AAP address as a parameter even if it isn't used at the time. `--tower-git-repository` is used to import the cluster data into AAP. TPA creates its own branch using `cluster_name` as the branch name, which allows you to use the same repository for all of your clusters. All other options to `tpaexec configure`, as described in [Cluster configuration](#), are still valid.

#### config.yml modification

`config.yml` includes the top-level dictionary `ansible_tower`, which causes `tpaexec provision` to treat the cluster as an AAP-enabled cluster.

Edit `config.yml` to ensure that `ansible_host` and `{private,public}_ip` are defined for each node and `ansible_host` is set to a value that AAP can resolve. Make any further changes or additions that you may need. See [Cluster configuration](#) for more details.

To generate inventory and other related files, run `tpaexec provision`.

## On the AAP UI

### Project

Add a project in AAP using the git repository as the source. Set the default EE of the project to use the TPA EE image.

**Allow Branch Override** is required when trying to use multiple inventories with a single project.

### Inventory

Add an empty inventory. Use the project as an external source to populate it using `inventory/00-cluster_name` as the inventory file.

We also recommend using **Overwrite local variables from remote inventory source** when not setting additional variables outside TPA's control in AAP.

### Credentials

Create a `vault` credential. You can retrieve the vault password using `tpaexec show-vault <cluster_dir>` on the TPA workstation.

To connect to your inventory nodes by way of SSH during deployment, make sure the machine credential is available in AAP.

### Template creation

To create a template:

1. Create a template that uses your project and your inventory.
2. Include these required credentials:
  - Vault credential
  - `EDB_SUBSCRIPTION_TOKEN` credential
  - Machine credential
3. Set two additional variables:

```
tpa_dir: /opt/EDB/TPA
cluster_dir: /runner/project
```

4. Select `deploy.yml` as the playbook.
5. To deploy your cluster, run a job based on the new template.

## 7.1.3 Use one project for multiple inventory

TPA uses a different branch name for each of your clusters in the associated git repository. This approach allows the use of a single project for multiple clusters.

### Set Allow branch override option

In the AAP project, enable the **Allow branch override** option.

## Define multiple inventories

TPA uses a different branch name for each of your clusters in the git repository. You can generate multiple inventories using the same project as the source by overriding the branch for each inventory.

## Define credentials per inventory

Ensure vault passwords are set accordingly per inventory since these differ on each TPA cluster.

### 7.1.4 Update TPA on AAP

Updating TPA on AAP involves some extra steps.

#### Update TPA workstation package

Update your TPA workstation package as any OS package depending on your OS. See [TPA installation](#) .

#### Use EE image with same version tag

Modify the EE image in AAP to use the same version tag as the workstation package version used.

#### Run tpaexec relink on your cluster directory

Ensure that any cluster using AAP is up to date by running `tpaexec relink <cluster_dir> --force` . An example of when you need to do this is after you have upgraded your TPA installation to a new version. Be sure to push any change committed by the `relink` command:

```
git status
```

```
On branch cluster_name
Your branch is ahead of tower/cluster_name by 1 commit.
  (use "git push" to publish your local commits)
...
```

```
git push tower
```

## Sync project and inventories

If they aren't set to use **Update revision on job launch** and **Update on launch**, sync the project in the AAP UI and related inventories, respectively.

### 7.1.5 Build an EE for TPA

#### Prerequisites

In order to build your own EE image, we recommend using `ansible-builder` .

You need:

1. `docker` or `podman`
2. `ansible-builder` and `ansible-navigator` python toolkits
3. TPA source code checked out at tag `vA.B.C` from [TPA repo](#) where `vA.B.C` is the TPA version you want to use.

## Environment file

`ansible-builder` uses an environment file to help generate a working EE image.

Here is a template example of such an environment file for TPA:

### execution-environment.yml

```
version: 3
images:
  base_image:
    name: registry.redhat.io/ansible-automation-platform-24/ee-minimal-rhel9:latest
dependencies:
  python: << TPA_REPO_CLONE_FOLDER >>/requirements-aap.txt
  galaxy: << TPA_REPO_CLONE_FOLDER >>/collections/requirements.yml
options:
  package_manager_path: /usr/bin/microdnf

additional_build_steps:
  append_final:
    - RUN mkdir -p /opt/EDB/TPA
    - COPY << TPA_REPO_CLONE_FOLDER >> /opt/EDB/TPA
    - ENV PYTHONPATH="${PYTHONPATH:+${PYTHONPATH}:}/opt/EDB/TPA/lib"
```

Different base image may require more `additional_build_steps`. See [ansible-builder](#) for advanced usage.

## EE build command

The following command should build the EE image for you:

```
ansible-builder build \
  --file=execution-environment.yml \
  --container-runtime=<docker/podman> \
  --tag=<your-registry>/<your-namespace>/tpa-ee:vA.B.C \
  --verbosity 2
```

## 7.2 TPA, Ansible, and sudo

TPA uses Ansible with `sudo` to execute tasks with elevated privileges on target instances. This page explains how Ansible uses `sudo` (which is in no way TPA-specific), and the consequences to systems managed with TPA.

TPA needs root privileges;

- to install packages (required packages using the operating system's native package manager, and optional packages using `pip`)
- to stop, reload and restart services (i.e Postgres, `repmgr`, `efm`, `etcd`, `haproxy`, `pgbouncer` etc.)
- to perform a variety of other tasks (e.g., gathering cluster facts, performing switchover, setting up cluster nodes)

TPA also needs to be able to use `sudo`. You can make it `ssh` in as root directly by setting `ansible_user: root`, but it will still use `sudo` to execute tasks as other users (e.g., `postgres`).

## 7.2.1 Ansible sudo invocations

When Ansible runs a task using `sudo`, you will see a process on the target instance that looks something like this:

```
/bin/bash -c sudo -H -S -n -u root /bin/bash -c \  
  ""echo BECOME-SUCCESS-kfoodiiprztsyerriqbjughhbemejgpc ; \  
  /usr/bin/python2"" && sleep 0
```

People who were expecting something like `sudo yum install -y xyzpkg` are often surprised by this. By and large, most tasks in Ansible will invoke a Python interpreter to execute Python code, rather than executing recognisable shell commands. (Playbooks may execute `raw` shell commands, but TPA uses such tasks only to bootstrap a Python interpreter.)

Ansible modules contain Python code of varying complexity, and an Ansible playbook is not just a shell script written in YAML format. There is no way to “extract” shell commands that would do the same thing as executing an arbitrary Ansible playbook.

There is one significant consequence of how Ansible uses `sudo`: [privilege escalation must be general](#). That, it is not possible to limit `sudo` invocations to specific commands in `sudoers.conf`, as some administrators are used to doing. Most tasks will just invoke `python`. You could have restricted `sudo` access to `python` if it were not for the random string in every command—but once Python is running as root, there’s no effective limit on what it can do anyway.

Executing Python modules on target hosts is just the way Ansible works. None of this is specific to TPA in any way, and these considerations would apply equally to any other Ansible playbook.

## 7.2.2 Recommendations

- Use SSH public key-based authentication to access target instances.
- Allow the SSH user to execute `sudo` commands without a password.
- Restrict access by time, rather than by command.

TPA needs access only when you are first setting up your cluster or running `tpaexec deploy` again to make configuration changes, e.g., during a maintenance window. Until then, you can disable its access entirely (a one-line change for both `ssh` and `sudo`).

During deployment, everything Ansible does is generally predictable based on what the playbooks are doing and what parameters you provide, and each action is visible in the system logs on the target instances, as well as the Ansible log on the machine where `tpaexec` itself runs.

Ansible’s focus is less to impose fine-grained restrictions on what actions may be executed and more to provide visibility into what it does as it executes, so elevated privileges are better assigned and managed by time rather than by scope.

## 7.2.3 SSH and sudo passwords

We *strongly* recommend setting up password-less SSH key authentication and password-less `sudo` access, but it is possible to use passwords too.

If you set `ANSIBLE_ASK_PASS=yes` and `ANSIBLE_BECOME_ASK_PASS=yes` in your environment before running `tpaexec`, Ansible will prompt you to enter a login password and a `sudo` password for the remote servers. It will then negotiate the login/`sudo` password prompt on the remote server and send the password you specify (which will make your playbooks take noticeably longer to run).

We do not recommend this mode of operation because we feel it is a more effective security control to completely disable access through a particular account when not needed than to use a combination of passwords to restrict access. Using public key authentication for `ssh` provides an effective control over who can access the server, and it’s easier to protect a single private key per authorised user than it is to protect a shared password or multiple shared passwords.

Also, if you limit access at the ssh/sudo level to when it is required, the passwords do not add any extra security during your maintenance window.

## 7.2.4 sudo options

To use Ansible with sudo, you must not set `requiretty` in `sudoers.conf`.

If needed, you can change the sudo options that Ansible uses (`-H -S -n`) by setting `become_flags` in the `[privilege_escalation]` section of `ansible.cfg`, or `ANSIBLE_BECOME_FLAGS` in the environment, or `ansible_become_flags` in the inventory. All three methods are equivalent, but please change the sudo options only if there is a specific need to do so. The defaults were chosen for good reasons. For example, removing `-S -n` will cause tasks to timeout if password-less sudo is incorrectly configured.

## 7.2.5 Logging

For playbook executions, the sudo logs will show mostly invocations of Python (just as it will show only an invocation of bash when someone uses `sudo -i`).

For more detail, the syslog will show the exact arguments to each module invocation on the target instance. For a higher-level view of why that module was invoked, the `ansible.log` on the controller shows what that task was trying to do, and the result.

If you want even more detail, or an independent source of audit data, you can run `auditd` on the server and use the SELinux log files. You can get still more fine-grained syscall-level information from `bpfftrace/bcc` (e.g., `opensnoop` shows every file opened on the system, and `execsnoop` shows every process executed on the system). You can do any or all of these things, depending on your needs, with the obvious caveat of increasing overhead with increased logging.

## 7.2.6 Local privileges

The *TPA installation*

mention sudo only as shorthand for “run these commands as root somehow”. Once TPA is installed and you have run `tpaexec setup`, TPA itself does not require elevated privileges on the local machine. (But if you use Docker, you must run `tpaexec` as a user that belongs to a group that is permitted to connect to the Docker daemon.)

## 7.3 TPA - PuTTY Configuration guide

In order to use PuTTY under Windows to connect via ssh to the AWS instances that were created by the TPA utility **tpaexec provision**, the keys will need to be converted from the private key format (`.pem`) generated by Amazon EC2 to the PuTTY format (`.ppk`).

```
# Provision the cluster

[tpa]$ tpaexec provision <clustername>
```

PuTTY has a tool named **PuTTYgen**, which can convert keys to the required format.

• - \*

### 7.3.1 Key conversion

#### Locate private key

Locate the private key in the cluster directory `<clustername>` - it will be named according to the `cluster_name` variable set in `config.yml` prefixed by `id_` - e.g. if the `cluster_name` is set to `testenv1`, then the private key will be called `id_testenv1`.

### Save key as .pem

Copy this file into your Windows filesystem & save it as a .pem file - in this example **id\_testenv1.pem** - cut and pasting into a text file will work fine for this.

### Key conversion

Start **PuTTYgen** and under Parameters, select appropriate Type of key to generate:

For older versions of **PuTTYgen**, select **SSH-2 RSA**; for recent versions select **RSA**

- Do not select SSH-1 (RSA)\*

Now choose **Load** - in the box that says **PuTTY Private Key Files (\*.ppk)** you will need to select **All Files (\*.\*)**

Select your **.pem** file and choose **Open**, then click **OK**.

Select **Save private key** and click **Yes** to ignore the warning about saving the key without a passphrase. Make sure that the file suffix is **.ppk** and choose the same name as for the **.pem** file; in this example the filename might be **id\_testenv1.ppk**

- - \*

## 7.3.2 Configure PuTTY

Start **PuTTY** and select **Session** from the **Category** window. In the **Host Name** panel, enter <user>@<IP address> and in the Port Panel, enter **22**

The <user> and <IP address> can be found in the <clustername>/ssh\_config file which gets created by the ``tpaexec provision`` utility.

In the Putty **Category** window, Select **Connection**, expand **SSH** and select **Auth**

For the panel marked **Private key file for authentication**, click **Browse** and select the .ppk file that was saved above, then select **Open**

In the Putty **Category** window, select **Session** again, enter a session name in **Saved Sessions**, and **Save**

You should now be able to connect to the AWS host via PuTTY by selecting this saved session.

- - \*

## 7.4 Troubleshooting

### 7.4.1 Recreate python virtual environment

Occasionally the python venv can get in an inconsistent state, in which case the easiest solution is to delete and recreate it. Symptoms of a broken venv can include errors during provisioning like:

```
TASK [Write Vagrantfile and firstboot.sh]_
  ↳ *****
failed: [localhost] (item=Vagrantfile) => {"changed": false, "checksum":
  ↳ "bf1403a17d897b68fa8137784d298d4da36fb7f9", "item": "Vagrantfile", "msg": "Aborting,
  ↳ target uses selinux but python bindings (libselinux-python) arent installed!"}
```

To create a new virtual environment (assuming tpaexec was installed into the default location):

```
sudo rm -rf /opt/EDB/TPA/tpa-venv
sudo /opt/EDB/TPA/bin/tpaexec setup
```

## 7.4.2 Strange AWS errors regarding credentials

If the time & date of the TPA server isn't correct, you can get AWS errors similar to this during provisioning:

```
TASK [Register key tpa_cluster in each region]
↳ *****
An exception occurred during task execution. To see the full traceback, use -vvv. The
↳ error was: ClientError: An error occurred (AuthFailure) when calling the
↳ DescribeKeyPairs operation: AWS was not able to validate the provided access
↳ credentials
failed: [localhost] (item=eu-central-1) => {"boto3_version": "1.8.8", "botocore_
↳ version": "1.11.8", "changed": false, "error": {"code": "AuthFailure", "message":
↳ "AWS was not able to validate the provided access credentials"}, "item": "eu-
↳ central-1", "msg": "error finding keypair: An error occurred (AuthFailure) when
↳ calling the DescribeKeyPairs operation: AWS was not able to validate the provided
↳ access credentials", "response_metadata": {"http_headers": {"date": "Thu, 27 Sep
↳ 2018 12:49:41 GMT", "server": "AmazonEC2", "transfer-encoding": "chunked"}, "http_
↳ status_code": 401, "request_id": "a0d905ba-188f-48fe-8e5a-c8d8799e3232", "retry_
↳ attempts": 0}}
```

Solution - set the time and date correctly.

```
sudo ntpdate pool.ntp.org
```

## 7.4.3 Logging

By default, all tpaexec logging will be saved in logfile <clusterdir>/ansible.log

To change the logfile location, set environment variable ANSIBLE\_LOG\_PATH to the desired location - e.g.

```
export ANSIBLE_LOG_PATH=~/.ansible.log
```

To increase the verbosity of logging, just add -v /-vv /-vvv /-vvvv /-vvvvv to tpaexec command line:

```
tpaexec deploy <clustername> -v
```

Where:

```
- v      shows the results of modules
- vv     shows the files from which tasks come
- vv     shows what commands are being executed on the target machines
- vv     enables connection debugging, what callbacks have been loaded
- vv     shows some additional ssh configuration, filepath information
```

## 7.4.4 Cluster test

An easy way to smoketest an existing cluster is to run:

```
tpaexec test <clustername>
```

This will do a functional test of the cluster components, followed by a performance test of the cluster, using pgbench. As pgbench can take a while to complete, benchmarking can be omitted by running:

```
tpaexec test <clustername> --excluded_tasks=pgbench
```

### 7.4.5 TPA server test

To check the installation of the TPA server itself, run:

```
tpaexec selftest
```

### 7.4.6 Including or excluding specific tasks

When re-running a `tpaexec` provision or deploy after a failure or when running tests, it can sometimes be useful to miss out tasks using TPA's *Selective task execution*.

## 7.5 Reference

### 7.5.1 Installation

#### Managing clusters in a disconnected or air-gapped environment

In a security controlled environment where no direct connection to the Internet is allowed, it is necessary to provide all packages needed by TPA to complete the deployment. This can be done via a local-repo on each node in the cluster. TPA supports the addition of custom repositories on each node via a *Creating and using a local repository* and the required packages can be downloaded using the `tpaexec download-packages` command.

#### Preparation

Choose an internet connected machine where you can install TPA and follow the instructions below to either copy an existing cluster configuration or create a new cluster.

#### Note

If the air-gapped server does not already have TPA installed, follow the instructions *Where to install TPA*

to install it.

If you have an existing cluster in a disconnected environment, all you need on the internet connected host is the `config.yml`. Create a directory and copy that file into it then run `tpaexec relink` on that directory to generate the remaining files that would normally be created by `tpaexec configure`.

Alternatively, to create a new configuration for an environment where the target instances will not have network access, configure a new cluster with this option:

```
tpaexec configure --use-local-repo-only ...
```

This will do everything that `--enable-local-repo` does, and disable the configuration for all other package repositories. On RedHat instances, this also includes disabling access to subscription-based services.

In an existing cluster, you can set `use_local_repo_only: yes` in `config.yml`:

```
cluster_vars:
  use_local_repo_only: yes
```

Note: that you do not need separate cluster configurations for internet connected and disconnected environments, the options below work in both.

More info on *Creating and using a local repository*

## Downloading packages

On the internet connected machine, ensure that you have *Installing Docker* and run:

```
tpaexec download-packages cluster-dir --os <OS> --os-version <version>
```

See detailed description for the *tpaexec download-packages* .

## Copying packages to the target environment

The resulting repository will be contained in the `cluster-dir/local-repo` directory. This is a complete package repo for the target OS. Copy this directory, from the connected controller to the disconnected controller that will be used to deploy the cluster. Place the directory in the same place, beneath the cluster directory. TPA will then copy packages to the instances automatically when `deploy` is run.

## Deploying in a disconnected environment

Ensure that the cluster `config.yml` has been configured as above in *Preparation* . Run `tpaexec provision` and `deploy` as you would normally.

## Updating in a disconnected environment

You can use the *Upgrading your cluster* command to perform updates in an air-gapped environment. Prior to running this command you must run `download-packages` on the connected controller and copy the updated repository to the disconnected controller.

## Distribution support

TPA detects and adapts to the distribution running on each target instance. This page lists platforms which are actively supported and ‘legacy distribution’ which have previously been supported. Deploying to a legacy platform is likely to work as long as you have access to the necessary packages, but this is not considered a supported use of TPA and is not suitable for production use.

Fully supported platforms are supported both as host systems for running TPA and target systems on which TPA deploys the Postgres cluster.

### Debian ARM64

- Debian 12/bookworm is fully supported

### Debian x86

- Debian 12/bookworm is fully supported
- Debian 11/bullseye is fully supported
- Debian 10/buster is a legacy distribution
- Debian 9/stretch is a legacy distribution
- Debian 8/jessie is a legacy distribution

## Ubuntu x86

- Ubuntu 24.04/noble is fully supported
- Ubuntu 22.04/jammy is fully supported
- Ubuntu 20.04/focal is a legacy distribution
- Ubuntu 18.04/bionic is a legacy distribution
- Ubuntu 16.04/xenial is a legacy distribution

## Oracle Linux x86

- Oracle Linux 9.x is fully supported (docker only)
- Oracle Linux 8.x is fully supported (docker only)
- Oracle Linux 7.x is a legacy distribution (docker only)

## RedHat x86

- RHEL/Rocky/AlmaLinux/Oracle Linux 9.x is fully supported (python3 only)
- RHEL/CentOS/Rocky/AlmaLinux 8.x is fully supported (python3 only)
- RHEL/CentOS 7.x is a legacy distribution (python2 only)

## RedHat ppc64le

- RHEL/Rocky/AlmaLinux 9.x is fully supported (python3 only)
- RHEL/AlmaLinux 8.x is fully supported (python3 only)

## SLES x86

- SLES 15.x is fully supported

## Platform-specific considerations

Some platforms may not work with the legacy distributions mentioned here. For example, Debian 8 and Ubuntu 16.04 are not available in *Installing Docker* .

## Running TPA in a Docker container

If you are using a system for which there are no *TPA installation* available, and it's difficult to run TPA after *Installing TPA from source* (for example, because it's not easy to obtain a working Python 3.12+ interpreter), your last resort may be to build a Docker image and run TPA inside a Docker container.

Please note that you do not need to run TPA in a Docker container in order to *Installing Docker* . It's always preferable to run TPA directly if you can (even on MacOS X).

## Quickstart

You must have Docker installed and working on your system already.

Run the following commands to clone the tpaexec source repository from Github and build a new Docker image named tpa/tpaexec :

```
git clone ssh://git@github.com/EnterpriseDB/tpa.git
cd tpa
docker build -f docker/Dockerfile --build-arg TPA_VER=$(git describe) -t_
↳ tpaexec:latest .
```

Double-check the created image:

```
docker image ls tpaexec
```

| REPOSITORY | TAG    | IMAGE ID     | CREATED        | SIZE  |
|------------|--------|--------------|----------------|-------|
| tpaexec    | latest | 3943dec4d660 | 20 minutes ago | 658MB |

```
docker run --rm tpaexec info
```

```
# TPAexec v23.38.0-38-g4dc030dc1

tpaexec=/usr/local/bin/tpaexec
TPA_DIR=/opt/EDB/TPA
PYTHON=/opt/EDB/TPA/tpa-venv/bin/python3 (v3.13.5, venv)
TPA_VENV=/opt/EDB/TPA/tpa-venv
ANSIBLE=/opt/EDB/TPA/tpa-venv/bin/ansible (v2.16.14)
```

Then you need to setup an alias for `tpaexec` on the shell session you are running:

```
alias tpaexec="docker run --rm -v $PWD:/work -v /var/run/docker.sock:/var/run/docker.
↳ sock tpaexec"
```

Now you can run commands like:

```
tpaexec configure cluster -a M1 --postgresql 15 --failover-manager patroni --platform_
↳ docker
tpaexec deploy cluster
```

## Installing Docker

Please consult the [Docker documentation](#) if you need help to install Docker and get started with it.

On MacOS X, you can [install "Docker Desktop forMac"](#)

and launch Docker from the application menu.

## TPA capabilities and supported software

- *Setting up the TPA Python environment*
- *Distribution support*

## Supported software

TPA can install and configure the following major components.

- Postgres
- EPAS (EDB Postgres Advanced Server)
- PGD 5, 4, 3.7

- pglogical 3, 2 (open source)
- pgd-cli and pgd-proxy
- HARP 2
- repmgr
- Barman
- pgbouncer
- haproxy (supported only for PGD 3.7)
- Failover Manager (EFM)
- Postgres Enterprise Manager (PEM)

## 7.5.2 Administration

### Reconciling changes made outside of TPA

Any changes made to a TPA created cluster that are not performed by changing the TPA configuration will not be saved in `config.yml`. This means that your cluster will have changes that the TPA configuration won't be able to recreate.

This page shows how configuration is managed with TPA and the preferred ways to make configuration changes. We then look at strategies to make, and reconcile, the results of making manual changes to the cluster.

### Why might I need to make manual configuration changes?

The most common scenario in which you may need to make configuration changes outside of TPA is if the operation you are performing is not supported by TPA. The two most common such operations are destructive changes, such as removing a node, and upgrading the major version of Postgres.

### Destructive changes

In general TPA will not remove previously deployed elements of a cluster, even if these are removed from `config.yml`. This sometimes surprises people because a strictly declarative system should always mutate the deployed artifacts to match the declaration. However, making destructive changes to production database can have serious consequences so it is something we have chosen not to support.

### Major-version Postgres upgrades

TPA does not yet provide an automated mechanism for performing major version upgrades of Postgres. Therefore if you need to perform an in-place upgrade on an existing cluster this must be performed using other tools such as `pg_upgrade` or `pgd node upgrade`.

### What can happen if changes are not reconciled?

A general issue with unreconciled changes is that if you deploy a new cluster using your existing `config.yml`, or provide your `config.yml` to EDB Support in order to reproduce a problem, it will not match the original cluster. In addition, there is potential for operational problems should you wish to use TPA to manage that cluster in future.

The operational impact of unreconciled changes varies depending on the nature of the changes. In particular whether the change is destructive, and whether the change blocks TPA from running by causing an error or invalidating the data in `config.yml`.

## Non-destructive, non-blocking changes

Additive changes are often accommodated with no immediate operational issues. Consider manually adding a user. The new user will continue to exist and cause no issues with TPA at all. You may prefer to manage the user through TPA in which case you can declare it in `config.yml` but the existence of a manually-added user will cause no operational issues.

Some manual additions can have more nuanced effects. Take the example of an extension which has been manually added. Because TPA does not make destructive changes, the extension will not be removed when `tpaexec deploy` is next run. **However**, if you made any changes to the Postgres configuration to accommodate the new extension these may be overwritten if you did not make them using one of TPA's supported mechanisms (see below).

Furthermore, TPA will not make any attempt to modify the `config.yml` file to reflect manual changes and the new extension will be omitted from `tpaexec upgrade` which could lead to incompatible software versions existing on the cluster.

## Destructive, non-blocking changes

Destructive changes that are easily detected and do not block TPA's operation will simply be undone when `tpaexec deploy` is next run. Consider manually removing an extension. From the perspective of TPA, this situation is indistinguishable from the user adding an extension to the `config.yml` file and running `deploy`. As such, TPA will add the extension such that the cluster and the `config.yml` are reconciled, albeit in the opposite way to that the user intended.

Similarly, changes made manually to configuration parameters will be undone unless they are:

1. Made in the `conf.d/9999-override.conf` file reserved for manual edits;
2. Made using `ALTER SYSTEM SQL`; or
3. Made `postgres_conf_settings` by adding `postgres_conf_settings`.

Other than the fact that option 3 is self-documenting and portable, there is no pressing operational reason to reconcile changes made by method 1 or 2.

## Destructive, blocking changes

Changes which create a more fundamental mismatch between `config.yml` can block TPA from performing operations. For example if you physically remove a node in a bare metal cluster, attempts by TPA to connect to that node will fail, meaning most TPA operations will exit with an error and you will be unable to manage the cluster with TPA until you reconcile this difference.

## How to reconcile configuration changes

In general, the reconciliation process involves modifying `config.yml` such that it describes the current state of the cluster and then running `tpaexec deploy`.

### Example: parting a PGD node

Deploy a minimal PGD cluster using the bare architecture and a configure command such as:

```
tpaexec configure mycluster \
  -a PGD-Always-ON \
  --platform bare \
  --edbpge 15 \
  --location-names a \
  --pgd-proxy-routing local
```

Part a node using this SQL, which can be executed from any node:

```
select * from bdr.part_node('node-2');
```

Rerun `deploy`. Note that, whilst no errors occur, the node is still parted. This can be verified using the command `pgd show-nodes` on any of the nodes. This is because TPA will not overwrite the metadata which tells PGD the node is parted.

!!!Note

It is not possible to reconcile the `config.yml` with this cluster state because TPA, and indeed PGD itself, has no mechanism to initiate a node in the 'parted' state. In principle you could continue to use TPA to continue this parted cluster, but this is not advisable. In most cases you will wish to continue to fully remove the node and reconcile `config.yml`.

The previous example parted a node from the PGD cluster, but left the node itself intact and still managed by TPA in a viable but unreconcilable state.

To completely decommission the node, it is safe to simply turn off the server corresponding to `node-2`. If you attempt to run `deploy` at this stage, it will fail early when it cannot reach the server.

To reconcile this change in `config.yml` simply delete the entry under `instances` corresponding to `node-2`. It will look something like this:

```
- Name: node-2
  public_ip: 44.201.93.236
  private_ip: 172.31.71.186
  location: a
  node: 2
  role:
  - bdr
  - pgd-proxy
  vars:
    bdr_child_group: a_subgroup
    bdr_node_options:
      route_priority: 100
```

You can now manage this node as usual using TPA. The original cluster still has metadata that refers to `node-2` as a node whose state is `PARTED`, which is not removed by default as it does not affect cluster functionality.

!!!Note

If you wish to join the original `node-2` back to the cluster after having removed it from `config.yml`, you can do so by restoring the deleted lines of `config.yml`, stopping Postgres, deleting the `PGDATA` directory on that node, and then repeating `tpaexec deploy`. As noted above, TPA will not remove an existing database, even if the corresponding entry is deleted from `config.yml`, so you need to perform this action manually.

TPA automatically generates a password for the superuser which you may view using `tpaexec show-password <cluster> <superuser-name>`. If you change the password manually (for example using the `/password` command in `psql`) you will find that after `tpaexec deploy` is next run, the password has reverted to the one set by TPA. To make the change through TPA, and therefore make it persist across runs of `tpaexec deploy`, you must use the command `tpaexec store-password <cluster> <superuser-name>` to specify the password, then run `tpaexec deploy`. This also applies to any other user created through TPA.

### Example: adding or removing an extension

A simple single-node cluster can be deployed with the following `config.yml`.

```

- ---
architecture: M1
cluster_name: singlenode

cluster_vars:
  postgres_flavour: postgresql
  postgres_version: 15
  preferred_python_version: python3

instance_defaults:
  image: tpa/debian:11
  platform: docker
  vars:
    ansible_user: root

instances:
- Name: nodeone
  node: 1
  role:
  - primary

```

You may manually add the `pgvector` extension by connecting to the node and running `apt install postgresql-15-pgvector` then executing the following SQL command: `CREATE EXTENSION vector;`. This will not cause any operational issues, beyond the fact that `config.yml` no longer describes the cluster as fully as it did previously. However, it is advisable to reconcile `config.yml` (or indeed simply use TPA to add the extension in the first place) by adding the following cluster variables.

```

cluster_vars:
  ...
  extra_postgres_packages:
    common:
      - postgresql-15-pgvector
  extra_postgres_extensions:
    - vector

```

After adding this configuration, you may manually remove the extension by executing the SQL command `DROP EXTENSION vector;` and then `apt remove postgresql-15-pgvector`. However if you run `tpaexec deploy` again without reconciling `config.yml`, the extension will be reinstalled. To reconcile `config.yml`, simply remove the lines added previously.

!!!Note

As noted previously, TPA will not honour destructive changes. So simply removing the lines from `config.yml` will not remove the extension. It is necessary to perform this operation manually then reconcile the change. !!!

## 7.5.3 Services

### Barman

When an instance is given the `barman` role in `config.yml`, TPA will configure it as a [Barman](#) server to take backups of any other instances that name it in their `backup` setting.

```
instances:

- Name: one
  backup: two
  ...

- Name: two
  role:
  - barman
  ...
```

Multiple `postgres` instances can have the same Barman server named as their `backup` ; equally, one `postgres` instance can have a list of Barman servers named as its `backup` and backups will be taken to all of the named servers.

The default Barman configuration will connect to PostgreSQL using `pg_receivewal` to take continuous backups of WAL, and will take a full backup of the instance using `rsync` over `ssh` twice weekly. Full backups and WAL are retained for long enough to enable recovery to any point in the last 4 weeks.

### Barman package version

By default, TPA installs the latest available version of Barman.

The version of the Barman package that is installed can be specified by including `barman_package_version: xxx` under the `cluster_vars` section of the `config.yml` file.

```
cluster_vars:
  ...
  barman_package_version: 1.56.2-1
  ...
```

You may use any version specifier that `apt` or `yum` would accept.

If your version does not match, try appending a `*` wildcard. This is often necessary when the package version has an epoch qualifier like `2:...` .

### Barman configuration

The Barman home directory on the Barman server can be set using the cluster variable `barman_home` ; its default value is `/var/lib/barman` .

On each Barman server, a global configuration file is created as `/etc/barman.conf` . This file contains default values for many Barman configuration variables. For each Postgres server being backed up, an additional Barman configuration file is created. For example, to back up the server `one` , the file is `/etc/barman.d/one.conf` , and the backups are stored in the subdirectory `one` in the Barman home directory. The configuration file and directory names can be changed from the backed-up instance's `backup_name` setting defined on the `vars` section before the provisioning step.

```
- Name: myPrimary
  backup: myBarman
  platform: bare
  ip_address: x.x.x.x
  node: 1
  role:
  - primary
```

(continues on next page)

(continued from previous page)

```
vars:
  backup_name: my_backup
```

The following variables can be set on the backed-up instance and are passed through into Barman's configuration with the prefix `barman_` removed:

| variable                                     | default                          |
|----------------------------------------------|----------------------------------|
| <code>barman_archiver</code>                 | false                            |
| <code>barman_log_file</code>                 | /var/log/barman.log              |
| <code>barman_backup_method</code>            | rsync                            |
| <code>barman_compression</code>              | pigz                             |
| <code>barman_reuse_backup</code>             | link                             |
| <code>barman_parallel_jobs</code>            | 1                                |
| <code>barman_backup_options</code>           | concurrent_backup                |
| <code>barman_immediate_checkpoint</code>     | false                            |
| <code>barman_network_compression</code>      | false                            |
| <code>barman_basebackup_retry_times</code>   | 3                                |
| <code>barman_basebackup_retry_sleep</code>   | 30                               |
| <code>barman_minimum_redundancy</code>       | 3                                |
| <code>barman_retention_policy</code>         | RECOVERY WINDOW OF 4 WEEKS       |
| <code>barman_last_backup_maximum_age</code>  | 1 WEEK                           |
| <code>barman_pre_archive_retry_script</code> |                                  |
| <code>barman_post_backup_retry_script</code> |                                  |
| <code>barman_post_backup_script</code>       |                                  |
| <code>barman_streaming_wals_directory</code> |                                  |
| <code>backup_name</code>                     | <i>backed up instance's name</i> |

## Backup scheduling

TPA installs a cron job in `/etc/cron.d/barman` which will run every minute and invoke `barman cron` to perform maintenance tasks.

For each instance being backed up, it installs another cron job in `/etc/cron.d/<backup_name>` which takes the backups of that instance. This job runs as determined by the `barman_backup_interval` variable for the instance, with the default being to take backups at 04:00 every Wednesday and Saturday.

## SSH keys

TPA will generate ssh key pairs for the `postgres` and `barman` users and install them into the respective `~/.ssh` directories, and add them to each other's `authorized_keys` file. The `postgres` user must be able to ssh to the `barman` server in order to archive WAL segments (if configured), and the `barman` user must be able to ssh to the Postgres instance to take or restore backups.

## barman and barman\_role Postgres users

TPA will create two Postgres users, `barman` and `barman_role`.

TPA versions <23.35 created the `barman` Postgres user as a `superuser`.

Beginning with 23.35 the `barman` user is created with `NOSUPERUSER`, so any re-deploys on existing clusters will remove the `superuser` attribute from the `barman` Postgres user. Instead, the `barman_role` is granted the required set of privileges and the `barman` user is granted `barman_role` membership.

This avoids granting the `superuser` attribute to the `barman` user, using the set of privileges provided in the [Barman Manual](#).

## Shared Barman server

### Note

To use the shared Barman functionality with clusters created using a TPA version earlier than 23.35, you must: a) upgrade to a version of TPA that supports creating shared Barman instances. b) after upgrading TPA, run `deploy` on `$first-cluster` so TPA can make necessary config changes for subsequent clusters to run smoothly against the shared Barman node.

Some deployments may want to share a single Barman server for multiple clusters. Shared Barman server deployment within `tpaexec` is supported via the `barman_shared` setting that can be set via `vars:` under the Barman server instance for the given cluster config that plans to use an existing Barman server. `barman_shared` is a boolean variable so possible values are `true` and `false` (default). When making any changes to the Barman config in a shared scenario, you must ensure that configurations across multiple clusters remain in sync so as to avoid a scenario where one cluster adds a specific configuration and a second cluster overrides it.

A typical workflow for using a shared Barman server across multiple clusters is described below.

1. Create a TPA cluster with an instance that has `barman` role (call it ‘first-cluster’ for this example).
2. In the second cluster (second-cluster for example), reference this particular Barman instance from `$clusters/first-cluster` as a shared Barman server instance and use `bare` as platform so we are not trying to create a new Barman instance when running provision. Also specify the IP address of the Barman instance that this cluster can use to access it.

```
- Name: myBarman
  node: 5
  role:
    - barman
  platform: bare
  ip_address: x.x.x.x
  vars:
    barman_shared: true
```

3. Once the second-cluster is provisioned but before running `deploy`, make sure that it can access the Barman server instance via `ssh`. You can allow this access by copying second-cluster’s public key to Barman server instance via `ssh-copy-id` and then do an `ssh` to make sure you can login without having to specify the password.

```
# add first-clusters key to the ssh-agent
cd $clusters/first-cluster
ssh-add id_first-clutser
cd $clusters/second-cluster
ssh-keyscan -t rsa,ecdsa -4 $barman-server-ip >> tpa_known_hosts
ssh-copy-id -i id_second-cluster.pub -o UserKnownHostsFile=tpa_known_hosts $user@
↪$barman-server-ip
ssh -F ssh_config $barman-server
```

4. Copy the Barman user’s keys from first-cluster to second-cluster

```
mkdir $clusters/second-cluster/keys
cp $clusters/first-cluster/keys/id_barman* clusters/second-cluster/keys
```

5. Run `tpaexec deploy $clusters/second-cluster`

### !!!Note Mixed-platform clusters

By declaring the shared Barman instance as `platform: bare` you might have changed your cluster to a mixed-platform cluster. This may require you to adjust other parts of `config.yml` to accommodate this change. Specifically, the `instance_defaults` section must only contain settings which are applicable to all instances in the cluster. If, for example, your `instance_defaults` contains a setting such as `type` which is only valid for `platform: aws` you must move that setting out of `instance_defaults` and into only the instances which use the AWS platform.

You must use caution when setting up clusters that share a Barman server instance. There are a number of important aspects you must consider before attempting such a setup.

1. Make sure that no two instances in any of the clusters sharing a Barman server use the same name. The `--cluster-prefixed-hostnames` option of `tpaexec configure` may be helpful in this respect.
2. Barman configuration and settings otherwise should remain in sync in all the clusters using a common Barman server to avoid a scenario where one cluster sets up a specific configuration and the others do not either because the configuration is missing or uses a different value.
3. The version of Postgres on instances being backed up across different clusters needs to be the same.
4. Different clusters using a common Barman server cannot specify different versions of Barman packages when attempting to override the default.

Some of these may be addressed in a future release as we continue to improve the shared Barman server support.

### !!!Warning

Be extremely careful when deprovisioning clusters sharing a common Barman node. Especially where the first cluster that deployed Barman uses non-bare platform. Deprovisioning the first cluster that originally provisioned and deployed Barman will effectively leave other clusters sharing the Barman node in an inconsistent state because the Barman node will already have been deprovisioned by the first cluster and it won't exist anymore.

When trying to upgrade to a specific package version, ensure the `barman_package_version` in `config.yml` is updated to reflect the desired version. The desired version can also be passed as an extra argument to the `tpaexec upgrade` command with:

```
- e barman_package_version="<desired version>"
```

Refer to the section on [package version selection and upgrade](#)

for more information.

To select Barman for upgrade, ensure the `--components` flag passed to the `tpaexec upgrade` command contains `barman` (or `all`)

Refer to the section on [component selection for upgrade](#)

for more information.

## EDB Postgres Distributed configuration

TPA can install and configure EDB Postgres Distributed (PGD), formerly known as BDR (Bi-directional replication) versions 3.7, 4.x, 5.x. and 6.x.

Access to PGD packages is through EDB's package repositories only. You must have a valid EDB subscription token to download the packages.

This documentation touches on several aspects of PGD configuration, but we refer you to the [PGDdocumentation](#) for an authoritative description of the details.

### Introduction

TPA will install PGD and any dependencies on all PGD instances along with Postgres itself.

After completing the basic Postgres setup and starting Postgres, TPA will then create the `bdr_database` and proceed to set up a PGD cluster through the various steps described below.

### Installation

TPA will install the correct PGD packages, depending on the version and flavour of Postgres in use (e.g., Postgres, Postgres Extended, or EPAS).

Set `bdr_version` to determine which major version of PGD to install (i.e., 3, 4, 5, 6). Set `bdr_package_version` to determine which exact package to install (e.g., '5.0\*' to install the latest 5.0.x).

### Overview of cluster setup

After installing the required packages, configuring Postgres to load PGD, and starting the server, TPA will go on to set up PGD nodes, groups, replication sets, and other resources.

Here's a summary of the steps TPA performs:

- Create a PGD node (using `bdr.create_node()`) for each participating instance
- Create one or more PGD node groups (using `bdr.create_node_group()`) depending on `bdr_node_groups`
- Create replication sets, if required, to control exactly which changes are replicated (depending on node group type and memberships, e.g., subscriber-only and witness nodes may need special handling)
- Join the relevant node groups on the individual instances
- Perform additional configuration, such as enabling subgroup RAFT or proxy routing.

(This process involves executing a complex sequence of queries, some on each instance in turn, and others in parallel. To make the steps easier to follow, TPA designates an arbitrary PGD primary instance as the “first\_bdr\_primary” for the cluster, and uses this instance to execute most of these queries. The instance is otherwise not special, and its identity is not significant to the PGD configuration itself.)

### Instance roles

Every instance with `bdr` in its `role` is a PGD instance, and implicitly also a `postgres` server instance.

A PGD instance with `readonly` in its role is a logical standby node (which joins the PGD node group with `pause_in_standby` set), eligible for promotion.

A PGD instance with `subscriber-only` in its role is a subscriber-only node, which receives replicated changes but does not publish them.

A PGD instance with `witness` in its role is a witness node.

Every PGD instance described above is implicitly also a `primary` instance. The exception is an instance with `replica` in its role; that indicates a physical streaming replica of an upstream PGD instance. Such instances are not included in any recommended PGD architecture, and not currently supported by TPA.

### Configuration settings

The settings mentioned below should ordinarily be set in `cluster_vars`, so that they are set uniformly for all the PGD instances in the cluster. You can set different values on different instances in some cases (e.g., `bdr_database`), but in other cases, the result is undefined (e.g., all instances must have exactly the same value of `bdr_node_groups`).

We strongly recommend defining your PGD configuration by setting uniform values for the whole cluster under `cluster_vars`.

### bdr\_database

The `bdr_database` (default: `bdrdb`) will be initialised with PGD.

### bdr\_client\_dsn\_attributes

Any *additional* parameter keywords supported by `libpq` can be included in `bdr_client_dsn_attributes`.

Do not include `host`, `port`, `dbname` and `user`, as these will already be included in the connection string.

!!!Note When `pgd-proxy` and `pgd-cli` are installed

Since `pgd-proxy` and `pgd-cli` are written in Go, they use Go drivers for connecting to Postgres.

These drivers do **not** support the full set of DSN attributes provided by the `libpq` C library.

If `pgd-proxy` and/or `pgd-cli` are installed and `bdr_client_dsn_attributes` includes parameters that are *unsupported* by the Go driver (such as `timeout`), two new variables must be included in the cluster configuration:

- `pgd_proxy_dsn_attributes`, which is used to create the connection strings in `pgd-proxy-conf`
- `pgd_cli_dsn_attributes`, which is used to create the connection strings in `pgd-cli-conf`

These two strings must **ONLY** contain parameter keywords compatible with the Go driver.

If the `bdr_client_dsn_attributes` does not include any unsupported parameters, this can be ignored and the `bdr_client_dsn_attributes` will be included in the connection strings for `pgd-proxy-conf` and `pgd-cli-conf`.

The setting of `bdr_node_group` (default: based on the cluster name) identifies which PGD cluster an instance should be a part of. It is also used to identify a particular cluster for external components (e.g., `pgd-proxy` or `harp-proxy`).

### bdr\_node\_groups

This is a list of PGD node groups that must be created before the group join stage (if the cluster requires additional subgroups).

In general, `tpaexec configure` will generate an appropriate value based on the selected architecture.

```
cluster_vars:
  bdr_node_groups:
    - name: topgroup
    - name: abc_subgroup
      node_group_type: data
      parent_group_name: topgroup
      options:
        location: abc
  ...
```

The first entry must be for the cluster's `bdr_node_group`.

Each subsequent entry in the list must specify a `parent_group_name`, and may specify the `node_group_type` (optional).

Each entry may also have an optional key/value mapping of group options. The available options vary by PGD version.

## **bdr\_child\_group**

If `bdr_child_group` is set for an instance (to the name of a group that is mentioned in `bdr_node_groups` ), it will join that group instead of `bdr_node_group` .

## **bdr\_commit\_scopes**

This is an optional list of [commit scopes](#) that must exist in the PGD database (available for PGD 4.1 and above).

```
cluster_vars:
  bdr_commit_scopes:
    - name: somescope
      origin: somegroup
      rule: ALL (somegroup) ON received ...`
    - name: otherscope
      origin: othergroup
      rule: ...
  ...
```

Each entry must specify the name of the commit scope, the name of the `origin` group, and the commit scope `rule` . The groups must correspond to entries in `bdr_node_groups` .

If you set `bdr_commit_scopes` explicitly, TPA will create, alter, or drop commit scopes as needed to ensure that the database matches the configuration. If you do not set it, it will leave existing commit scopes alone.

## **preferred\_first\_bdr\_primary**

This is the name of the instance that TPA will use for PGD-related tasks that only run on a single node. This includes executing DDL/DML, which is then propagated to the other PGD nodes by replication. It also includes reading data for initializing a new PGD node. If not set, or set to an instance which is not usable, TPA will choose an instance arbitrarily from the suitable candidates if any.

The instance selected by TPA is stored in the fact `first_bdr_primary` . The list of viable instances is stored in the fact `first_bdr_primary_candidates` . Neither of these facts can be set from `config.yml` , they are noted here for use in hooks or debugging.

!!!Note

The choice of `preferred_first_bdr_primary` does not influence which nodes are elected as write leaders in the deployed cluster.

If you wish to set `preferred_first_bdr_primary` for a single operation, you can pass it using the Ansible `-e` option rather than specifying it in `config.yml` . For example:

```
tpaexec deploy . -e preferred_first_bdr_primary=instance_name
```

This can be useful if you are adding or rebuilding a PGD node and wish to ensure that data is clone from a node in the same data centre for example.

Similarly, you can set `preferred_first_bdr_primary` as an instance variable, with different instances on each node. For example, this configuration would mean that if `node1` was rebuilt, the `tpaexec deploy` command would prefer `node2` as the source of data with which to rebuild it.

```
instances:
  - Name: node1
```

(continues on next page)

(continued from previous page)

```
...
instance_vars:
  preferred_first_bdr_primary: node2
```

If you use this configuration, it should be applied after the node is initially deployed; having multiple values of `preferred_first_bdr_primary` during a first deploy is not recommended.

## Miscellaneous notes

### Hooks

TPA invokes the `bdr-node-pre-creation`, `bdr-post-group-creation`, and `bdr-pre-group-join` *TPA hooks* during the PGD cluster setup process.

### Database collations

TPA checks that the PGD database on every instance in a cluster has the same collation (`LC_COLLATE`) setting. Having different collations in databases in the same PGD cluster is a data loss risk.

### Older versions of PGD

TPA no longer actively supports or tests the deployment of BDR v1 (with a patched version of Postgres 9.4), v2 (with Postgres 9.6), or any PGD versions below v3.7.

### Configuring EFM

TPA will install and configure EFM when `failover_manager` is set to `efm`.

The version of EFM to be used can be configured by changing the `efm_version`:

```
cluster_vars:
...
  failover_manager: efm
  efm_version: 5.2
...
```

It can also be defined when executing `tpaexec configure`:

```
tpaexec configure my-cluster-dir -a M1 --failover_manager efm --efm-version 5.2
```

Note that EFM is only available via EDB's package repositories and requires a valid subscription.

### EFM configuration

TPA will generate `efm.nodes` and `efm.properties` with the appropriate instance-specific settings, with remaining settings set to the respective default values. TPA will also place an `efm.notification.sh` script which basically contains nothing by default and leaves it up to the user to fill it in however they want. TPA will override the default settings for `auto.allow.hosts` and `stable.nodes.file` to simplify adding agents to the cluster.

See the [EFM documentation](#)

for more details on EFM configuration.

## efm\_user\_password\_encryption

Must be either `scram-sha-256` or `md5`

Set `efm_user_password_encryption` to control the auth-method for the efm Postgres user's auth-method in `pg_hba.conf` as well as the algorithm used when generating its encrypted password.

### Note

The default value for `efm_user_password_encryption` is `scram-sha-256` for new clusters.

```
efm_user_password_encryption: scram-sha-256 # or can be set to `md5`
```

## efm\_conf\_settings

You can use `efm_conf_settings` to set specific parameters. These must be written as entries in an Ansible dictionary, in `key: value` form

See the *documentation on the `efm.properties` file* <[https://www.enterprisedb.com/docs/efm/latest/04\\_configuring\\_efm/01\\_cluster\\_properties/](https://www.enterprisedb.com/docs/efm/latest/04_configuring_efm/01_cluster_properties/)> for

details on which settings can be configured.

```
cluster_vars:
  efm_conf_settings:
    notification.level: WARNING
    ping.server.ip: <well known address in network>
```

If you make changes to values under `efm_conf_settings`, TPA will always restart EFM to activate the changes.

## EFM witness

TPA will install and configure EFM as witness on instances whose `role` contains `efm-witness` (when using EFM, you can also use the role `witness`, as it gets implicitly converted to `efm-witness`). For such instances the `upstream` property must be specified to point to the designated primary database instance.

## Repmgr

EFM works as a failover manager and therefore TPA will still install `repmgr` for setting up postgresql replicas on postgres versions 11 and below. `repmgrd` i.e. `repmgr`'s daemon remains disabled in this case and `repmgr`'s only job is to provided replication setup functionality.

For postgres versions 12 and above, any cluster that uses EFM will use `pg_basebackup` to create standby nodes and not use `repmgr` in any form.

## Node Promotability

TPA determines whether a node is eligible for promotion by EFM during failover based on the node's role and replication topology. The following rules are applied when generating the EFM configuration:

- **Witness nodes** (`efm-witness` role) are never promotable.
- **Nodes with the `efm-not-promotable` role** are not eligible for promotion. This can be used to prevent specific standbys, such as DR or reporting nodes, from being promoted to primary.
- **Cascading standbys** (nodes that are not directly replicating from the primary) are also not promotable.

- All other nodes are considered promotable by default.

To explicitly prevent a standby from being promoted, add `efm-not-promotable` to the node's `roles` list in your cluster configuration. This ensures that EFM will not attempt to promote this node during failover events.

### Minor update for EFM using `tpaexec upgrade`

!!!Note EFM is *not* a package update

Due to how EFM versions are released, each EFM minor version has its own package (`edb-efm49` for v4.9, `edb-efm50` for v5.0, etc.) which installs its own distinct binary and configuration directories.

TPA installs the package for the new version and copies over the configuration files from the existing EFM version's configuration directory into the new EFM version's configuration directory, and removes the existing EFM version's package, service file and binary directory for cleanup.

Because EFM upgrade depends on a previous version being installed and configured (the source), TPA first checks to ensure this is true. In the event a source EFM version is not installed (it's binary directory does not exist) or not configured (it's configuration directory does not exist), the upgrade will exit with an error. When running the upgrade, if the target EFM version (specified by `efm_version` in the `config.yml`) already has a binary directory and configuration directory, upgrade skips over EFM as it is already installed and configured to the desired version.

!!!Note

EFM upgrade is supported from EFM 4.9 to 5.0 even if not a minor upgrade per say, no additional steps are required for this upgrade to happen.

Refer to the section on *component selection for upgrade*

for more information.

!!!Note

EFM can't work in a mixed version environment. Due to this limitation, there is no good reason to support a partial upgrade of the EFM component. TPA will **always** upgrade the EFM software on the complete set of EFM nodes when `efm` is part of the list of components to upgrade. TPA will ignore `update_hosts` when deciding which nodes should have EFM upgraded. !!!

### Configuring haproxy

TPA will install and configure haproxy on instances whose `role` contains `haproxy`.

By default, haproxy listens on `127.0.0.1:5432` for requests forwarded by *Repmgr redirect pgbouncer* running on the same instance. You must specify a list of `haproxy_backend_servers` to forward requests to.

TPA will install the latest available version of haproxy by default. You can install a specific version instead by setting `haproxy_package_version: 1.9.15*` (for example).

Note: see limitations of using wildcards in `package_version` in *tpaexec-configure*.

Haproxy packages are selected according to the type of architecture. An EDB managed haproxy package may be used but requires a subscription. Packages from PGDG extras repo can be installed if required.

You can set the following variables on any `haproxy` instance.

| Variable                       | Default value                      | Description                                                                                                                                                                                                          |
|--------------------------------|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>haproxy_bind_address</i>    | 127.0.0.1*                         | The address to which haproxy should bind <i>0.0.0.0</i> if <i>failover_manager</i> is <i>patroni</i> . * Users should change this value to something more restrictive and appropriate for their cluster networking** |
| <i>haproxy_port</i>            | 5432 (5444 for EPAS)               | The TCP port haproxy should listen on                                                                                                                                                                                |
| <i>haproxy_read_only_port</i>  | 5433 (5445 for EPAS)               | TCP port for read-only load-balancer                                                                                                                                                                                 |
| <i>haproxy_backend_servers</i> | None                               | A list of Postgres instance names                                                                                                                                                                                    |
| <i>haproxy_maxconn</i>         | <code>`max_connections`×0.9</code> | The maximum number of connections allowed per backend server; the default is derived from the backend's <i>max_connections</i> setting                                                                               |
| <i>haproxy_peer_enabled</i>    | True*                              | Add known haproxy hosts as <i>peer</i> list. *`False` if <i>failover_manager</i> is <i>harp</i> or <i>patroni</i> .                                                                                                  |

## Read-Only load-balancer

Haproxy can be configured to listen on an additional port for read-only access to the database. At the moment this is only supported with the Patroni failover manager. The backend health check determines which postgres instances are currently acting as replicas and will send traffic using a roundrobin load balancing algorithm.

The read-only load balancer is disabled by default but can be turned on using the `cluster_vars` variable `haproxy_read_only_load_balancer_enabled`.

## Patroni Health Checks

When Patroni is used as the failover manager, HAProxy performs health checks against the Patroni REST API. TPA automatically configures these checks based on Patroni's security settings:

- **HTTP:** If `patroni_ssl_enabled` is `false`, standard HTTP checks (option `httpchk`) are used against the Patroni API port (default 8008).
- **HTTPS (Basic Auth or no Auth):** If `patroni_ssl_enabled` is `true` and `patroni_authentication_mode` is `basic`, HAProxy uses `check-ssl` and is configured to trust the cluster's CA certificate to validate the Patroni API server certificate (using `verify required ca-file <path>`).
- **HTTPS (mTLS):** If `patroni_authentication_mode` is `mtls`, HAProxy uses `check-ssl` and is configured with the cluster's CA certificate *and* a unique client certificate/key pair for authentication (using `verify required ca-file <path> crt <path>`). TPA automatically generates and deploys these client certificates for HAProxy.

## Server options

TPA will generate `/etc/haproxy/haproxy.cfg` with a backend that has a `default-server` line and one line per backend server. All but the first one will be marked as "backup" servers.

Set `haproxy_default_server_extra_options` to a list of options on the haproxy instance to add options to the `default-server` line; and set `haproxy_server_options` to a list of options on the backend server to add options (which will override the defaults) to the individual server lines for each backend. TPA automatically adds necessary `check` and `check-ssl` options based on the failover manager configuration.

The size of the RSA key size used for `ssl-dh-param-file` can be modified by adding the variable `ha_proxy_dh_params_key_size` in the `cluster_vars` section.

## Example

```
cluster_vars:
  ha_proxy_dhparams_key_size: 4096
instances:
- Name: one
  vars:
    haproxy_server_options:
      - maxconn 33
- Name: two
...
- Name: proxy
  role:
    - haproxy
  vars:
    haproxy_backend_servers:
      - one
      - two
    haproxy_default_server_extra_options:
      - on-error mark-down
      - on-marked-down shutdown-sessions
```

## Configuring HARP

TPA will install and configure HARP when `failover_manager` is set to `harp`, which is the default for BDR-Always-ON clusters.

## Installing HARP

You must provide the `harp-manager` and `harp-proxy` packages. Please contact EDB to obtain access to these packages.

## Variables for HARP configuration

See the [HARP documentation](#)

for more details on HARP configuration.

| Variable                                 | Default value         | Description                                                                          |
|------------------------------------------|-----------------------|--------------------------------------------------------------------------------------|
| <code>cluster_name</code>                | <code>''</code>       | The name of the cluster.                                                             |
| <code>harp_consensus_protocol</code>     | <code>''</code>       | The consensus layer to use ( <i>etcd</i> or <i>bdr</i> )                             |
| <code>harp_location</code>               | <code>location</code> | The location of this instance (defaults to the <i>location</i> parameter)            |
| <code>harp_ready_status_duration</code>  | <code>10</code>       | Amount of time in seconds the node's readiness status will persist if not refreshed. |
| <code>harp_leader_lease_duration</code>  | <code>6</code>        | Amount of time in seconds the Lead Master lease will persist if not refreshed.       |
| <code>harp_lease_refresh_interval</code> | <code>2000</code>     | Amount of time in milliseconds between refreshes of the Lead Master lease.           |

continues on next page

Table 3 – continued from previous page

| Variable                               | Default value   | Description                                                                                                                                                                                 |
|----------------------------------------|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>harp_dcs_reconnect_interval</i>     | <i>1000</i>     | The interval, measured in ms, between attempts that a disconnected node tries to reconnect to the DCS.                                                                                      |
| <i>harp_dcs_priority</i>               | <i>500</i>      | In the case two nodes have an equal amount of lag and other qualified criteria to take the Lead Master lease, this acts as an additional ranking value to prioritize one node over another. |
| <i>harp_stop_database_when_fenced</i>  | <i>false</i>    | Rather than simply removing a node from all possible routing, stop the database on a node when it is fenced.                                                                                |
| <i>harp_fenced_node_on_dcs_failure</i> | <i>false</i>    | If HARP is unable to reach the DCS then fence the node.                                                                                                                                     |
| <i>harp_maximum_lag</i>                | <i>1048576</i>  | Highest allowable variance (in bytes) between last recorded LSN of previous Lead Master and this node before being allowed to take the Lead Master lock.                                    |
| <i>harp_maximum_camo_lag</i>           | <i>1048576</i>  | Highest allowable variance (in bytes) between last received LSN and applied LSN between this node and its CAMO partner(s).                                                                  |
| <i>harp_camo_enforcement</i>           | <i>lag_only</i> | Whether CAMO queue state should be strictly enforced.                                                                                                                                       |
| <i>harp_use_unix_sock</i>              | <i>false</i>    | Use unix domain socket for manager database access.                                                                                                                                         |
| <i>harp_request_timeout</i>            | <i>250</i>      | Time in milliseconds to allow a query to the DCS to succeed.                                                                                                                                |
| <i>harp_watch_poll_interval</i>        | <i>500</i>      | Milliseconds to sleep between polling DCS. Only applies when <i>harp_consensus_protocol</i> is <i>bdr</i> .                                                                                 |
| <i>harp_proxy_timeout</i>              | <i>1</i>        | Builtin proxy connection timeout, in seconds, to Lead Master.                                                                                                                               |
| <i>harp_proxy_keepalive</i>            | <i>5</i>        | Amount of time builtin proxy will wait on an idle connection to the Lead Master before sending a keepalive ping.                                                                            |
| <i>harp_proxy_max_client_conn</i>      | <i>75</i>       | Maximum number of client connections accepted by harp-proxy ( <i>max_client_conn</i> )                                                                                                      |
| <i>harp_ssl_password_command</i>       | <i>None</i>     | a custom command that should receive the obfuscated sslpassword in the stdin and provide the handled sslpassword via stdout.                                                                |
| <i>harp_db_request_timeout</i>         | <i>10s</i>      | similar to dcs -> request_timeout, but for connection to the database itself.                                                                                                               |
| <i>harp_local_etcd_only</i>            | <i>None</i>     | limit harp manager endpoints list to only contain the local etcd node instead of all etcd nodes                                                                                             |

You can use the *harp-config hook*

to execute tasks after the HARP configuration files have been installed (e.g., to install additional configuration files).

## Consensus layer

The `--harp-consensus-protocol` argument to `tpaexec configure` is mandatory for the BDR-Always-ON architecture.

## etcd

If the `--harp-consensus-protocol etcd` option is given to `tpaexec configure`, then TPA will set `harp_consensus_protocol` to `etcd` in `config.yml` and give the `etcd` role to a suitable subset of the instances, depending on your chosen layout.

HARP v2 requires etcd v3.5.0 or above, which is available in the `products/harp/release` package repositories provided by EDB.

You can configure the following parameters for etcd:

| Variable                      | Default value | Description                                  |
|-------------------------------|---------------|----------------------------------------------|
| <code>etcd_peer_port</code>   | 2380          | The port used by etcd for peer communication |
| <code>etcd_client_port</code> | 2379          | The port used by clients to connect to etcd  |

## bdr

If the `--harp-consensus-protocol bdr` option is given to `tpaexec configure`, then TPA will set `harp_consensus_protocol` to `bdr` in `config.yml`. In this case the existing PGD instances will be used for consensus, and no further configuration is required.

## Configuring a separate user for harp proxy

If you want harp proxy to use a separate readonly user, you can specify that by setting `harp_dcs_user: username` under `cluster_vars`. TPA will use `harp_dcs_user` setting to create a readonly user and set it up in the DCS configuration.

## Configuring a separate user for harp manager

If you want harp manager to use a separate user, you can specify that by setting `harp_manager_user: username` under `cluster_vars`. TPAexec will use that setting to create a new user and grant it the `bdr_superuser` role.

## Custom SSL password command

The command provided by `harp_ssl_password_command` will be used by HARP to de-obfuscate the `sslpassword` given in connection string. If `sslpassword` is not present then `harp_ssl_password_command` is ignored. If `sslpassword` is not obfuscated then `harp_ssl_password_command` is not required and should not be specified.

## Configuring the harp service

You can configure the following parameters for the harp service:

| Variable                                     | Default value | Description                                                                                                                                                                                       |
|----------------------------------------------|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>harp_manager_restart_on_failure</code> | <i>false</i>  | If <i>true</i> , the <i>harp-manager</i> service is overridden so it's restarted on failure. The default is <i>false</i> to comply with the service installed by the <i>harp-manager</i> package. |

## Configuring harp http(s) health probes

You can enable and configure the http(s) service for harp that will provide api endpoints to monitor service's health.

| Variable                 | Default value                                                                                                                                                                                                                                                                                                                      | Description                                                                          |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| <i>harp_http_options</i> | <pre>&lt;br /&gt;`enable: false`&lt;br /&gt;`se- cure:      false`&lt;br /&gt;`host:    &lt;inven- tory_hostname&gt;`&lt;br /&gt;`port:    8080`&lt;br /&gt;`probes:&lt;br /&gt;          `timeout: 10s`&lt;br /&gt;`endpoint: "host=&lt;proxy_name&gt; port=&lt;6432&gt;      db- name=&lt;bdrdb&gt; user=&lt;username&gt;"</pre> | Configure the http section of harp config.yml that defines the http(s) api settings. |

The variable can contain these keys:

```
enable: false
secure: false
cert_file: "/etc/tpa/harp_proxy/harp_proxy.crt"
key_file: "/etc/tpa/harp_proxy/harp_proxy.key"
host: <inventory_hostname>
port: 8080
probes:
  timeout: 10s
endpoint: "<valid dsn>"
```

The `cert_file` and `key_file` keys are both required if you use `secure: true` and are willing to use your own certificate and key.

You must ensure that both certificate and key are available at the given location on the target node before running `deploy`.

Leave both `cert_file` and `key_file` empty if you want TPA to generate a certificate and key for you using a cluster specific CA certificate. TPA CA certificate won't be 'well known', you will need to add this certificate to the trust store of each machine that will probe the endpoints. The CA certificate can be found on the cluster directory on the TPA node at: `<cluster_dir>/ssl/CA.crt` after `deploy`.

see harp documentation for more information on the available api endpoints.

## Configuring Postgres Enterprise Manager (PEM)

TPA will install and configure PEM when `tpaexec configure` command is run with `--enable-pem` command line option.

The default behavior with `--enable-pem` is to enable `pem-agent` role for all `postgres` instances in the cluster. `pem-agent` role will also be added to `barman` nodes when `--enable-pg-backup-api` command line option is used alongside `--enable-pem`.

A dedicated instance named `pemserver` will also be added to the cluster.

Since PEM server uses postgres backend; `pemserver` instance implicitly uses `postgres` role as well which ensures that `pemserver` gets a valid postgres cluster configured for use as PEM backend. All configuration options available for a normal postgres instance are valid for PEM's backend postgres instance as well. See following for details:

- `pg_hba.conf`
- `postgresql.conf`

Note that PEM is only available via EDB's package repositories and therefore requires a valid subscription.

## Supported architectures

PEM is supported with all architectures via the `--enable-pem` configuration command line option, with the exception of the BDR-Always-ON architecture when used with EDB Postgres Extended. You can optionally edit the generated cluster config (`config.yml`) and assign or remove `pem-agent` role from any postgres instance in the cluster in order to enable or disable PEM there.

## PEM component package versions

By default, TPA installs the latest available version of PEM agent and PEM server.

The version of the PEM agent and PEM server packages that are installed can be specified by including `pem_agent_package_version: xxx` and `pem_server_package_version: xxx` under the `cluster_vars` section of the `config.yml` file.

```
cluster_vars:
  ...
  pem_agent_package_version: 9.7.0-1.el9
  pem_server_package_version: 9.7.0-1.el9
  ...
```

You may use any version specifier that apt or yum would accept.

If your version does not match, try appending a `*` wildcard. This is often necessary when the package version has an epoch qualifier like `2:...`.

## PEM configuration

TPA will configure pem agents and pem server with the appropriate instance-specific settings, with remaining settings set to the respective default values. Some of the configuration options may be exposed for user configuration at some point in future.

!!!Note PEM components are upgraded independently.

Passing `pem-server` to `--components` upgrades ONLY the PEM server, whereas passing `pem-agent` upgrades ONLY the PEM agents. If upgrading separately, it is recommended to upgrade the agents before the server. When both are upgraded together, TPA upgrades the agents before the server.

## Passing additional options when registering PEM agents

TPA registers each PEM agent in the cluster using the `pemworker` utility's `--register agent` command.

A list of additional registration options can be passed by including `pemagent_registration_opts` in the cluster config.

For example:

```
pemagent_registration_opts:
- --enable-smtp true
- --enable-heartbeat-connection
- --allow-batch-probes true
- -l DEBUG1
```

The [PEM documentation](#)

lists more information about registration options.

### Useful extensions for the nodes with pem agent

By default, TPA will add `sql_profiler`, `edb_wait_states` and `query_advisor` extensions to any instances that have pem-agent role.

This list of default extensions for pem-agent nodes can be overridden by setting `pemagent_extensions` in `config.yml`. If this list is empty, no extensions will be automatically included.

### Providing an external certificate for PEM server SSL authentication

By default, the PEM server creates a self-signed certificate pair, `server-pem.crt` and `server-pem.key` and configures the webserver to use them for HTTPS access.

The size of `server-pem.key` can be modified adding the variable `pem_rsa_key_size` to the `cluster_vars` section:

```
(...)
cluster_vars:
  pem_rsa_key_size: 4096
```

The size of the CA certificate expedited by the PEM database can also be modified adding the variable `pem_db_ca_certificate_key_size` to the `cluster_vars` section:

```
(...)
cluster_vars:
  pem_db_ca_certificate_key_size: 4096
```

By default, the self-signed TLS certificate has a period of 10 years. The period can be adjusted by defining `openssl_certificate_validity` in the `cluster_vars` section. If the certificate needs to be replaced, use the `pem_web_server_renew_tls_certificates` variable on `tpaexec deploy:tpaexec deploy (...) -e pem_web_server_renew_tls_certificates=true`.

To provide your own certificate pair, create a directory under the root of the cluster directory named `ssl/pemserver` and place the certificate pair inside.

```
cluster directory
├─ ssl
│   └─ pemserver
│       ├── externally-provided.crt
│       └─ externally-provided.key
```

Next, set the variables `pem_server_ssl_certificate` and `pem_server_ssl_key` with the respective file names as values for the `vars:` under the pem server instance or `cluster_vars` in the cluster config file.

TPA will handle copying these files over to the pem server instance and configure the webserver accordingly.

```

- Name: pemserver
  location: main
  node: 4
  role:
  - pem-server
  vars:
    pem_server_ssl_certificate: externally-provided.crt
    pem_server_ssl_key: externally-provided.key

```

## Organising instances in PEM

PEM supports collating servers and/or agents into **groups** and **clusters**. You can specify this grouping from TPA by setting the following instance variables.

On instances with the `pem-agent` role, but without `pem-server` role:

- `monitoring_group` specifies the name of the group to which the Postgres server on this instance will be assigned. If not specified it defaults to the value of `pem_server_group`.
- `monitoring_agent_group` specifies the name of the group to which the PEM agent on this instance will be assigned. If not specified it defaults to the value of `pem_agent_group`.
- `monitoring_cluster` specifies the name of the cluster to which the Postgres server on this instance will be assigned. If not specified, no cluster will be assigned/created.
- `monitoring_agent_cluster` specifies the name of the cluster to which the PEM agent on this instance will be assigned. If not specified, no cluster will be assigned/created.

!!!Note

PEM only permits a given cluster name to exist in a single group, so some combinations of these values are not viable. For example if you specify the same `monitoring_cluster` and `monitoring_agent_cluster`, but different `monitoring_group` and `monitoring_agent_group`, this implies the same cluster name would appear in two different groups. TPA allows PEM to handle such inconsistencies rather than attempting to prevent them. This generally means deploy will succeed but you may not get exactly the organisation of servers and agents you expected.

- `pem_server_group` specifies the name of the default group to which Postgres servers registered with this PEM Server will be assigned. If not specified it defaults to `PEM Server Directory`.
- `pem_agent_group` specifies the name of the default group to which PEM Agents registered with this PEM Server will be assigned. If not specified it defaults to `PEM Agents`.

Note that the Postgres server and PEM Agent on the instance with `pem-server` role are always assigned to `pem_server_group` and `pem_agent_group` respectively. TPA does not support adding this server or agent to a cluster.

## Shared PEM server

Some deployments may want to use a single PEM server for monitoring and managing multiple clusters in the organization. Shared pem server deployment within `tpaexec` is supported via the `pem_shared` variable that you could set via `vars`: under the pem server instance for the given cluster config that plans to use an existing pem server. `pem_shared` is a boolean variable so possible values are `true` and `false`(default). When declaring a pemserver instance as shared, we tell the given cluster config that pemserver instance is in fact managed by a separate cluster config that provisioned and deployed the pem server in the first place. So any changes we wanted to make to the pem server instance including postgres backend for pem would be managed by the cluster where pemserver instance is NOT declared as a shared pem instance.

A typical workflow for using a shared pem server across multiple clusters would look something like this:

1. Create a tpaexec cluster with a single instance that has `pem-server` role (call it 'pem-cluster' for this example). We could as easily use the same workflow in a scenario where pem is provisioned as part of a larger cluster and not just a single instance that runs as pemserver but we use a single node cluster because it is easier to use that as an example and arguably easy to maintain as well.
2. In the other cluster (pg-cluster for example), reference this particular pemserver from `$clusters/pem-cluster` as a shared pem server instance and use `bare` as platform so we are not trying to create a new pemserver instance. Also specify the IP address of the pemserver that this cluster can use to access pemserver instance.

```
- Name: pemserver
  node: 5
  role:
  - pem-server
  platform: bare
  public_ip: 13.213.53.205
  private_ip: 10.33.15.102
  vars:
    pem_shared: true
```

3. Before running deploy in the postgres cluster, make sure that pg-cluster can access pem server instance via ssh. You can allow this access by copying pg-cluster's public key to pem server instance via `ssh-copy-id` and then do an ssh to make sure you can login without having to specify the password.

```
# add pem-clusters key to the ssh-agent (handy for `aws` platform)
$ cd $clusters/pem-cluster
$ ssh-add id_pem-clutser
$ cd $clusters/pg-cluster
$ ssh-keyscan -4 $pem-server-ip >> known_hosts
$ ssh-copy-id -i id_pg-cluster.pub -o UserKnownHostsFile=tpa_known_hosts $user@$pem-
→server-ip
$ ssh -F ssh_config pemserver
```

4. Update postgresql config on pem server node so it allows connections from the new pg-cluster. You can modify existing `pg_hba.conf` on pem server by adding new entries to `pem_postgres_extra_hba_settings` under vars: in pem-cluster's config.yml. For example:

```
instances:
- Name: pemserver
  location: main
  node: 1
  role:
  - pem-server
  vars:
    pem_postgres_extra_hba_settings:
      - "# Allow pem connections from pg-cluster1.quire"
      - hostssl pem +pem_agent 10.33.15.108/32 cert
      - "# Allow pem connections from pg-cluster1.upside"
      - hostssl pem +pem_agent 10.33.15.104/32 cert
      - "# Allow pem connections from pg-cluster2.zippy"
      - hostssl pem +pem_agent 10.33.15.110/32 cert
      - "# Allow pem connections from pg-cluster2.utopic"
      - hostssl pem +pem_agent 10.33.15.109/32 cert
```

and then run `tpaexec provision $clusters/pem-cluster` followed by `tpaexec deploy $clusters/pem-cluster`. When complete, nodes from your new pg-cluster should be able to speak with pem server backend.

5. In order to make sure pem agents from the nodes in pg-cluster can connect and register with the pem server backend, you must first `export EDB_PEM_CREDENTIALS_FILE=/path/to/pem/credentials/file` before you run `tpaexec deploy`. Credentials file is a text file that contains your access credentials to the pem-server's backend postgres instance in the `username:password` format.

```
cat pem_creds
```

```
postgres:f1I%fw!QmWevdzw#EL#$Ulu1cWhg7&RT
```

If you don't know the backend password, you can get that by using `show-password tpaexec` command.

```
tpaexec show-password $pem-clusterdir $user
```

6. Run `tpaexec deploy $clusters/pg-cluster` so pem is deployed on the new pg-cluster while using shared pem server instance.

!!!Note Mixed-platform clusters

By declaring the shared PEM instance as `platform: bare` you might have changed your cluster to a mixed-platform cluster. This may require you to adjust other parts of `config.yml` to accommodate this change. Specifically, the `instance_defaults` section must only contain settings which are applicable to all instances in the cluster. If, for example, your `instance_defaults` contains a setting such as `type` which is only valid for `platform: aws` you must move that setting out of `instance_defaults` and into only the instances which use the AWS platform.

PEM UI runs on https interface so you can connect with a running instance of PEM server via <https://protect\TU\textdollarpem-server-ip/pem>. Login credentials for PEM UI are set to the postgres backend user which uses `postgres` or `enterprisedb` for `postgresql` and `epas` flavours respectively. `tpaexec's show-password` command will show the password for the backend user. For example:

```
tpaexec show-password $clusterdir $user
```

See [PEM documentation](#)

for more details on PEM configuration and usage.

### Minor update for PEM using `tpaexec upgrade`

!!!Note PEM major version upgrades from version 9 to version 10.

The PEM server package includes an SQL upgrade script which handles upgrading the database according to the PEM schema version. Upgrading from PEM v9.8.0 to PEM 10.1.1 using `tpaexec upgrade` has been tested and works, but is a major version upgrade.

Passing `pem-server` to `--components` upgrades ONLY the PEM server, whereas passing `pem-agent` upgrades ONLY the PEM agents. If upgrading separately, it is recommended to upgrade the agents before the server. When both are upgraded together, TPA upgrades the agents before the server.

When trying to upgrade the PEM server to a specific package version, ensure the `pem_server_package_version` in `config.yml` is updated to reflect the desired version.

The desired version can also be passed as an extra argument to the `tpaexec upgrade` command with:

```
tpaexec upgrade <cluster_dir> \
-e pem_agent_package_version="<desired version>" \
-e pem_server_package_version="<desired version>" \
--components=pem-agent,pem-server
```

Refer to the section on [package version selection and upgrade](#) for more information.

To select PEM agents for upgrade, ensure the `--components` flag passed to the `tpaexec upgrade` command contains `pem-agent` (or `all`)

To select the PEM server for upgrade, ensure the `--components` flag passed to the `tpaexec upgrade` command contains `pem-server` (or `all`)

Refer to the section on [component selection for upgrade](#)

for more information.

## PgBouncer

### PgBouncer package version

By default, TPA installs the latest available version of PgBouncer.

The version of the PgBouncer package that is installed can be specified by including `pgbouncer_package_version: xxx` under the `cluster_vars` section of the `config.yml` file.

```
cluster_vars:
  ...
  pgbouncer_package_version: 1.8*
  ...
```

You may use any version specifier that apt or yum would accept.

If your version does not match, try appending a `*` wildcard. This is often necessary when the package version has an epoch qualifier like `2:...`

### Configuring PgBouncer

TPA will install and configure PgBouncer on instances whose `role` contains `pgbouncer`.

By default, PgBouncer listens for connections on port 6432 and, if no `pgbouncer_backend` is specified, forwards connections to `127.0.0.1:5432` (which may be either Postgres or [Configuring haproxy](#), depending on the architecture).

!!!Note Using PgBouncer to route traffic to the primary

If you are using the M1 architecture with `repmgr` you can set `repmgr_redirect_pgbouncer: true` hash under `cluster_vars` to have PgBouncer connections directed to the primary. The PgBouncer will be automatically updated on failover to route to the new primary. You should use this option in combination with setting `pgbouncer_backend` to the primary instance name to ensure that the cluster is initially deployed with PgBouncer configured to route to the primary.

| Variable                               | Default value                      | Description                                                                                                                           |
|----------------------------------------|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <code>pgbouncer_port</code>            | 6432                               | The TCP port pgbouncer should listen on                                                                                               |
| <code>pgbouncer_backend</code>         | 127.0.0.1                          | A Postgres server to connect to                                                                                                       |
| <code>pgbouncer_backend_port</code>    | 5432                               | The port that the <code>pgbouncer_backend</code> listens on                                                                           |
| <code>pgbouncer_max_client_conn</code> | <code>`max_connections`×0.9</code> | The maximum number of connections allowed; the default is derived from the backend's <code>max_connections</code> setting if possible |
| <code>pgbouncer_auth_user</code>       | <code>pgbouncer_auth_user</code>   | Postgres user to use for authentication                                                                                               |

## Databases

By default, TPA will generate `/etc/pgbouncer/pgbouncer.databases.ini` with a single wildcard `*` entry under `[databases]` to forward all connections to the backend server. You can set `pgbouncer_databases` as shown in the example below to change the database configuration.

## Authentication

PgBouncer will connect to Postgres as the `pgbouncer_auth_user` and execute the (already configured) `auth_query` to authenticate users.

The `pgbouncer_get_auth()` function used as the `auth_query` by PgBouncer is created in a single database, the `pgbouncer_auth_database`. Execute permissions are granted on this function to the `pgbouncer_auth_user`.

## Example

```
instances:

- Name: one
  vars:
    max_connections: 300

- Name: two

- Name: proxy
  role:
    - pgbouncer
  vars:
    pgbouncer_backend: one
    pgbouncer_databases:
      - name: dbname
        options:
          pool_mode: transaction
          dbname: otherdb
      - name: bdrdb
        options:
          host: two
          port: 6543
```

## Minor update for PgBouncer using `tpaexec upgrade`

When trying to upgrade to a specific package version, ensure the `pgbouncer_package_version` in `config.yml` is updated to reflect the desired version.

The desired version can also be passed as an extra argument to the `tpaexec upgrade` command with:

```
tpaexec upgrade <cluster_dir> \
  -e pgbouncer_package_version="<desired version>" \
  --components=pgbouncer
```

Refer to the section on *package version selection and upgrade* for more information.

To select PgBouncer for upgrade, ensure the `--components` flag passed to the `tpaexec upgrade` command contains `pgbouncer` (or `all`)

Refer to the section on *component selection for upgrade*

for more information.

## pgd-proxy

### pgd-proxy package version

By default, TPA installs the latest available version of `pgd-proxy`.

The version of the `pgd-proxy` package that is installed can be specified by including `pgd_proxy_package_version: xxx` under the `cluster_vars` section of the `config.yml` file.

```
cluster_vars:
  ...
  pgd_proxy_package_version: 5.0.0-1
  ...
```

You may use any version specifier that `apt` or `yum` would accept.

If your version does not match, try appending a `*` wildcard. This is often necessary when the package version has an epoch qualifier like `2: . . .`.

### Configuring pgd-proxy

TPA will install and configure `pgd-proxy` for the PGD-Always-ON architecture with PGD 5 on any instance with `pgd-proxy` in its role.

(By default, the *PGD-Always-ON*

will run `pgd-proxy` on all the data nodes in every location, but you can instead create any number of additional proxy instances with `--add-proxy-nodes-per-location 3`.)

### Configuration

`pgd-proxy` is configured at PGD level via SQL functions.

| Hash                           | Function                                   | Description                                                                |
|--------------------------------|--------------------------------------------|----------------------------------------------------------------------------|
| <code>pgd_proxy_options</code> | <code>bdr.alter_proxy_option()</code>      | <code>pgd-proxy</code> configuration, e.g. port                            |
| <code>bdr_node_groups</code>   | <code>bdr.alter_node_group_option()</code> | configuration for the proxy's node group, e.g. <i>enable_proxy_routing</i> |
| <code>bdr_node_options</code>  | <code>bdr.alter_node_option()</code>       | routing configuration for individual PGD nodes                             |

See the PGD documentation for more details.

You can use the *pgd-proxy-config hook*

to execute tasks after the PGD PROXY configuration files have been installed (e.g., to install additional configuration files).

### bdr\_node\_groups

Group-level options related to `pgd-proxy` can be set under `bdr_node_groups` along with other node group options:

```
cluster_vars:
  bdr_node_groups:
    - name: group1
```

(continues on next page)

(continued from previous page)

```
options:
  enable_proxy_routing: true
```

Note that `enable_proxy_routing` must be explicitly set to `true` for `pgd-proxy` to be enabled for the group.

### bdr\_node\_options

Node-level options related to `pgd-proxy` can be set under `bdr_node_options` on any PGD instance:

```
instances:

- Name: first
  vars:
    bdr_node_options:
      route_priority: 42
```

### pgd\_proxy\_options

Options for a `pgd-proxy` instance itself, rather than the group or nodes it is attached to, can be set under `default_pgd_proxy_options` under `cluster_vars` (which applies to all proxies), or under `pgd_proxy_options` on any `pgd-proxy` instance:

```
cluster_vars:
  default_pgd_proxy_options:
    listen_port: 6432
    read_listen_port: 6433

instances:

- Name: someproxy
  vars:
    pgd_proxy_options:
      listen_port: 9000
      read_listen_port: 9001
```

In this case, while other instances will get their `listen_port` setting from `cluster_vars`, `someproxy` overrides that default setting and configures its own `listen_port` in the instances' `vars` section.

### PGD proxy http(s) health probes

You can enable and configure the `http(s)` service for PGD proxy that will provide api endpoints to monitor the proxy's health.

`pgd_http_options` under `cluster_vars` or instance `vars` will store all the settings that defines the `http(s)` api which live under the `http` subsection of the `proxy` top section of `pgd-proxy-config.yml`.

The variable can contain these keys:

```
enable: false
secure: false
cert_file: "/etc/tpa/harp_proxy/harp_proxy.crt"
key_file: "/etc/tpa/harp_proxy/harp_proxy.key"
host: <inventory_hostname>
```

(continues on next page)

(continued from previous page)

```
port: 8080
probes:
  timeout: 10s
endpoint: "<valid dsn>"
```

The `cert_file` and `key_file` keys are both required if you use `secure: true` and are willing to use your own certificate and key.

You must ensure that both certificate and key are available at the given location on the target node before running `deploy`.

Leave both `cert_file` and `key_file` empty if you want TPA to generate a certificate and key for you using a cluster specific CA certificate. TPA CA certificate won't be 'well known', you will need to add this certificate to the trust store of each machine that will probe the endpoints. The CA certificate can be found on the cluster directory on the TPA node at: `<cluster_dir>/ssl/CA.crt` after `deploy`.

see `pgd-proxy` documentation for more information on the available api endpoints.

### Updating pgd-proxy using `tpaexec upgrade`

When trying to upgrade to a specific package version, ensure the `pgd_proxy_package_version` in `config.yml` is updated to reflect the desired version. This version should be kept in line with the `bdr_package_version` and `pgdcli_package_version` versions specified for the cluster. Defining `bdr_package_version` and omitting `pgd_proxy_package_version` and `pgdcli_package_version` from `config.yml` will ensure the same value is used across all three components.

The desired version can also be passed as an extra argument to the `tpaexec upgrade` command with:

```
tpaexec upgrade <cluster_dir> \
  -e pgd_proxy_package_version="<desired version>" \
  --components=pgd-proxy
```

Refer to the section on *package version selection and upgrade* for more information.

To select `pgd-proxy` for upgrade, ensure the `--components` flag passed to the `tpaexec upgrade` command contains `pgd-proxy` (or `all`)

Refer to the section on *component selection for upgrade*

for more information.

## pglogical

### pglogical package version

By default, TPA installs the latest available version of `pglogical`.

The version of the `pglogical` package that is installed can be specified by including `pglogical_package_version: xxx` under the `cluster_vars` section of the `config.yml` file.

```
cluster_vars:
  ...
  pglogical_package_version: 2.2.0*
  ...
```

You may use any version specifier that `apt` or `yum` would accept.

If your version does not match, try appending a `*` wildcard. This is often necessary when the package version has an epoch qualifier like `2:...`.

## pglogical configuration

TPA can configure pglogical replication sets (publications) and subscriptions with pglogical v2 and pglogical v3.

```
instances:

- node: 1
  Name: kazoo
  ...
  vars:
    publications:
      - type: pglogical
        database: example
        name: some_publication_name
        replication_sets:
          - name: custom_replication_set
            ...

- node: 2
  Name: keeper
  vars:
    subscriptions:
      - type: pglogical
        database: example
        name: some_subscription_name
        publication:
          name: some_publication_name
        replication_sets:
          - default
          - default_insert_only
          - custom_replication_set
            ...
```

The pglogical extension will be created by default if you define publications or subscriptions with `type: pglogical`, but it is up to you to determine which version will be installed (e.g., subscribe to the `products/pglogical3/` release repository for pglogical3).

## Introduction

TPA can configure everything needed to replicate changes between instances using pglogical, and can also alter the replication setup based on `config.yml` changes.

To publish changes, you define an entry with `type: pglogical` in `publications`. To subscribe to these changes, you define an entry with `type: pglogical` in `subscriptions`, as shown above.

Pglogical does not have a named publication entity (in the sense that built-in logical replication has `CREATE PUBLICATION`). A publication in `config.yml` just assigns a name to a collection of replication sets, and subscriptions can use this name to refer to the desired provider.

To use pglogical replication, both publishers and subscribers need a named local pglogical node. TPA will create this node with `pglogical.create_node()` if it does not exist. For publications, the publication name is used as the

pglogical node name. There can be only one pglogical node in any given database, so you can have only one entry in `publications` per database.

However, pglogical subscriptions *do* have a name of their own. TPA will create subscriptions with the given `name`, and use a default value for the pglogical node name based on the instance's name and the name of the database in which the subscription is created. You can specify a different `node_name` if required—for example, when you have configured a publication in the same database, so that all subscriptions in that database must share the same pglogical node.

TPA does some basic validation of the configuration—it will point out the error if you spell `replication_sets` as `replciation_sets`, or try to subscribe to a publication that is not defined, but it is your responsibility to specify a meaningful set of publications and subscriptions.

TPA will configure pglogical after creating users, extensions, and databases, but before any PGD configuration. You can set *Creating Postgres users* and *Creating Postgres databases* to create databases for replication, and use the *postgres-config-final*

hook to populate the databases before pglogical is configured.

## Publications

An entry in `publications` must specify a `name` and `database`, and may specify a list of named `replication_sets` with optional attributes, as well as a list of table or sequence names.

```
publications:
- type: pglogical
  database: example
  name: some_publication_name
  replication_sets:
  - name: default
    replicate_insert: true
    replicate_update: true
    replicate_delete: true
    replicate_truncate: true
    autoadd_tables: false
    autoadd_sequences: false
    autoadd_existing: true
  - name: custom_replication_set
    tables:
    - name: sometable
    - name: "some-schema".othertable
      columns: [a, b, c]
      row_filter: a > 42
      synchronize_data: true
    sequences:
    - name: someseq
      synchronize_data: true
    - name: "some-schema".otherseq
```

Each replication set may specify optional attributes such as `replicate_insert` and `autoadd_existing`. If specified, they will be included as named parameters in the call to `pglogical.create_replication_set()`, otherwise they will be left out and the replication set will be created with pglogical's defaults instead.

Apart from manipulating the list of relations belonging to the replication set using the `autoadd_*` parameters in `pglogical3`, you can also explicitly specify a list of tables or sequences. The name of each relation may be schema-qualified

(unqualified names are assumed to be in `public`), and the entry may include optional attributes such as `row_filter` (for tables only) or `synchronize_data`, as shown above.

## Subscriptions

An entry in `subscriptions` must specify a `name` and `database`, define a publication to subscribe to, and may specify other optional attributes of the subscription.

```
subscriptions:
- type: pglogical
  database: example
  name: some_subscription_name
  node_name: optional_pglogical_node_name
  publication:
    name: some_publication_name
  # Optional attributes:
  synchronize_structure: true
  synchronize_data: true
  forward_origins: [all]
  strip_origins: false
  apply_delay: 1 second
  writer: heap
  writer_options:
    - magic
    - key=value
    - just-a-string
  # Optional attributes that can be changed for an existing
  # subscription:
  replication_sets:
    - default
    - default_insert_only
    - custom_replication_set
  enabled: true
```

A subscription can set `publication.name` (as shown above) to define which publication to subscribe to. If there is more than one publication with that name (across the entire cluster), you may specify the name of an instance to disambiguate. If you want to refer to publications by name, don't create multiple publications with the same name on the same instance.

```
- type: pglogical
  ...
  publication:
    name: some_publication_name
    instance: kazoo

  # OR

  provider_dsn: "host=... dbname=..."
```

Instead of referring to publications by name, you may explicitly specify a `provider_dsn` instead. In this case, the given DSN is passed to `pglogical.create_subscription()` directly (and `publication` is ignored). You can use this mechanism to subscribe to instances outside the TPA cluster.

The other attributes in the example above are optional. If defined, they will be included as named parameters in the

call to `pglogical.create_subscription()`, otherwise they will be left out. (Some attributes shown are specific to `pglogical3`.)

## Configuration changes

For publications, you can add or remove replication sets, change the attributes of a replication set, or change its membership (the tables and sequences it contains).

If you change `replicate_*` or `autoadd_*`, TPA will call `pglogical.alter_replication_set()` accordingly (but note that you cannot change `autoadd_existing` for existing replication sets, and the `autoadd_*` parameters are all `pglogical3`-specific).

If you change the list of tables or sequences for a replication set, TPA will reconcile these changes by calling `pglogical.alter_replication_set_{add,remove}_{table,sequence}()` as needed.

However, if you change `synchronize_data` or other attributes for a relation (table or sequence) that is already a member of a replication set, TPA will not propagate the changes (e.g., by dropping the table and re-adding it with a different configuration).

For subscriptions, you can only change the list of `replication_sets` and enable or disable the subscription (`enabled: false`).

In both cases, any replication sets that exist but are not mentioned in the configuration will be removed (with `pglogical.alter_subscription_remove_replication_set()` on the subscriber, or `pglogical.drop_replication_set()` on the publisher—but the default replication sets named `default`, `default_insert_only`, and `ddl_sql` will not be dropped.)

If you edit `config.yml`, remember to run `tpaexec provision` before running `tpaexec deploy`.

## Interaction with PGD

It is possible to use PGD and `pglogical` together in the same database if you exercise caution.

PGD v3 uses `pglogical3` internally, and will create a `pglogical` node if one does not exist. There can be only one `pglogical` node per database, so if you configure a `pglogical` publication in `bdr_database`, the instance's `bdr_node_name` must be the same as the publication's name. Otherwise, the node will be created for the publication first, and `bdr.create_node()` will fail later with an error about a node name conflict. Any subscriptions in `bdr_database` must use the same `node_name` too.

## Limitations

- There is currently no support for `pglogical.replication_set_{add,remove}_ddl()`
- There is currently no support for `pglogical.replication_set_add_all_{tables,sequences}()`
- There is currently no support for `pglogical.alter_subscription_{interface,writer_options}()` or `pglogical.alter_subscription_{add,remove}_log()`
- `pglogical v1` support is not presently tested.

## repmgr

TPA will install `repmgr` on all postgres instances that have the `failover_manager` instance variable set to `repmgr`; this is the default setting.

The directory of the `repmgr` configuration file defaults to `/etc/repmgr/<version>`, where `<version>` is the major version of postgres being installed on this instance, but can be changed by setting the `repmgr_conf_dir` variable for the instance. The configuration file itself is always called `repmgr.conf`.

The default `repmgr` configuration will set up automatic failover between instances configured with the role `primary` and the role `replica`.

### repmgr package version

By default, TPA installs the latest available version of `repmgr`.

The version of the `repmgr` package that is installed can be specified by including `repmgr_package_version: xxx` under the `cluster_vars` section of the `config.yml` file.

```
cluster_vars:
  ...
  repmgr_package_version: 4.0.5-1.pgdg90+1
  ...
```

You may use any version specifier that `apt` or `yum` would accept.

If your version does not match, try appending a `*` wildcard. This is often necessary when the package version has an epoch qualifier like `2:...`.

### repmgr configuration

The following instance variables can be set:

```
repmgr_priority: sets priority in the config file
repmgr_location: sets location in the config file
repmgr_reconnect_attempts: sets reconnect_attempts in the config file, default 6
repmgr_reconnect_interval: sets reconnect_interval in the config file, default 10
repmgr_use_slots: sets use_replication_slots in the config file, default 1
repmgr_failover: sets failover in the config file, default automatic
```

Any extra settings in `repmgr_conf_settings` will also be passed through into the `repmgr` config file.

### repmgr on PGD instances

On PGD instances, `repmgr_failover` will be set to `manual` by default.

### Minor update for repmgr using tpaexec upgrade

!!!Warning `repmgr` minor-versioning is DIFFERENT: the THIRD digit (X.Y.Z) is the minor version

A minor release upgrade involves updating `repmgr` from one minor release to another minor release within the same major release (e.g. 5.3.1 to 5.3.2). An upgrade between minor releases of differing major releases (e.g. 5.2.1 to 5.3.2) is a MAJOR upgrade. TPA does NOT support major version upgrades of `repmgr`.

The desired version can also be passed as an extra argument to the `tpaexec upgrade` command with:

```
tpaexec upgrade <cluster_dir> \
  -e repmgr_package_version="<desired version>" \
  --components=repmgr
```

Refer to the section on *package version selection and upgrade* for more information.

To select `repmgr` for upgrade, ensure the `--components` flag passed to the `tpaexec upgrade` command contains `repmgr` (or `all`)

Refer to the section on *component selection for upgrade*

for more information.

## 7.5.4 Configuration

### Configuring EDB Repos 2.0 repositories

TPA sources EDB software from [EDB Repos 2.0](#) . To use EDB Repos 2.0 you must export `EDB_SUBSCRIPTION_TOKEN=xxx` before you run `tpaexec`. You can get your subscription token from [the web interface](#) .

!!!Note

If you create your `config.yml` file using the `tpaexec configure` command, the `edb_repositories` key will be automatically populated with the necessary repositories for your selected configuration, so you shouldn't need to edit it.

```
cluster_vars:
  edb_repositories:
    - enterprise
    - postgres_distributed
```

This example will install the 'enterprise' subscription repository as well as 'postgres\_distributed' giving access to EPAS and PGD products. On Debian or Ubuntu systems, it will use the APT repository and on RedHat or SLES systems, it will use the rpm repositories, through the yum or zypper frontends respectively.

### Configuring APT repositories

This page explains how to configure APT package repositories on Debian and Ubuntu systems.

You can define named repositories in `apt_repositories` , and decide which ones to use by listing the names in `apt_repository_list` :

```
cluster_vars:
  apt_repositories:
    Example:
      key_id: XXXXXXXXX
      key_url: https://repo.example.com/path/to/XXXXXXXXX.asc
      repo: >-
        deb https://repo.example.com/repos/Example/ xxx-Example main

  apt_repository_list:
    - PGDG
    - Example
```

This configuration would install the GPG key (with id `key_id` , obtained from `key_url` ) and a new entry under `/etc/apt/sources.list.d` with the given `repo` line (or lines) for the PGDG repository (which is already defined by default) and the new Example repository.

When you configure additional repositories, remember to include PGDG in `apt_repository_list` if you still want to install PGDG packages.

You can set `apt_repository_list: []` to not install any repositories.

### Configuring YUM repositories

This page explains how to configure YUM package repositories on RedHat systems.

You can define named repositories in `yum_repositories` , and decide which ones to use by listing the names in `yum_repository_list` :

```

cluster_vars:
  yum_repositories:
    Example:
      rpm_url: >-
        https://repo.example.com/repos/Example/example-repo.rpm

    Other:
      description: "Optional repository description"
      baseurl: https://other.example.com/repos/Other/$basearch
      gpgkey:
        https://other.example.com/repos/Other/gpg.XXXXXXXXXXXXXXXXXX.key

  yum_repository_list:
    - EPEL
    - PGDG
    - Example
    - Other

```

This example shows two ways to define a YUM repository.

If the repository has a “repo RPM” (a package that customarily installs the necessary `/etc/yum.repos.d/*.repo` file and any GPG keys needed to verify signed packages from the repository), you can just point to it.

Otherwise, you can specify a description, a `baseurl`, and a `gpgkey` URL, and TPA will create a `/etc/yum.repos.d/Other.repo` file for you based on this information.

The EPEL and PGDG repositories are defined by default. The EPEL repository is required for correct operation, so you must always include EPEL in `yum_repository_list`. You should also include PGDG if you want to install PGDG packages.

You can set `yum_repository_list: []` to not install any repositories (but things will break without an alternative source of EPEL packages).

If you need to perform any special steps to configure repository access, you can use a *TPA hooks* to create the `.repo` file yourself:

```

- name: Define Example repository
  copy:
    dest: /etc/yum.repos.d/example.repo
    owner: root
    group: root
    mode: "0644"
    content: |
      [example]
      name=Example repo
      baseurl=https://repo.example.com/repos/Example/
      enabled=1
      gpgkey=https://repo.example.com/repokey.asc
      gpgcheck=1

```

In this case, you do not need to list the repository in `yum_repository_list`.

## Creating and using a local repository

If you create a local repository within your cluster directory, TPA will make any packages in the repository available to cluster instances. This is an easy way to ship extra packages to your cluster.

Optionally, you can also instruct TPA to configure the instances to use *only* this repository, i.e., disable all others. In this case, you must provide *all* packages required during the deployment, starting from basic dependencies like rsync, Python, and so on.

You can create a local repository manually, or have TPA create one for you. Instructions for both are included below.

### Note

Specific instructions are available for *Managing clusters in a disconnected or air-gapped environment* environment.

## Creating a local repository with TPA

TPA includes tools to help create such a local repository. Specifically the `--enable-local-repo` switch can be used with `tpaexec configure` to create an empty directory structure to be used as a local repository, and `tpaexec download-packages` populates that structure with the necessary packages.

## Creating the directory structure

To configure a cluster with a local repository, run:

```
tpaexec configure --enable-local-repo ...
```

This will generate your cluster configuration and create a `local-repo` directory and OS-specific subdirectories. See below for *Local repo layout*.

## Populate the repository and generate metadata

Run `tpaexec download-packages` to download all the packages required by a cluster into the `local-repo`. The resulting repository will contain the full dependency tree of all packages so the entire cluster can be installed from this repository. Metadata for the repository will also be created automatically meaning it is ready to use immediately.

## Creating a local repository manually

### Local repo layout

To create a local repository manually, you must first create an appropriate directory structure. When using `--enable-local-repo`, TPA will create a `local-repo` directory and OS-specific subdirectories within it (e.g., `local-repo/Debian/12`), based on the OS you select for the cluster. We recommend that this structure is also used for manually created repositories.

For example, a cluster running RedHat 8 might have the following layout:

```
local-repo/
|-- RedHat
|   |-- 8.5 -> 8
|   |-- 8
|       |-- repodata
```

For each instance, TPA will look for the following subdirectories of `local-repo` in order and use the first one it finds:

- `<distribution>/<version>`, e.g., `RedHat/8.5`

- `<distribution>/<major version>`, e.g., `RedHat/8`
- `<distribution>/<release name>`, e.g., `Ubuntu/focal`
- `<distribution>`, e.g., `Debian`
- The `local-repo` directory itself.

If none of these directories exists, of course, TPA will not try to set up any local repository on target instances.

### Populating the repository and generating metadata

The steps detailed below must be completed before running `tpaexec deploy`.

To populate the repository, copy the packages you wish to include into the appropriate directory. Then generate metadata using the correct tool for your system as detailed below.

#### **Note**

You must generate the metadata on the control node, i.e., the machine where you run `tpaexec`. TPA will copy the metadata and packages to target instances.

#### **Note**

You must generate the metadata in the subdirectory that the instance will use, i.e., if you copy packages into `local-repo/Debian/12`, you must create the metadata in that directory, not in `local-repo/Debian`.

### Debian/Ubuntu repository metadata

For Debian-based distributions, install the `dpkg-dev` package:

```
sudo apt-get update && sudo apt-get install -y dpkg-dev
```

Now you can use `dpkg-scanpackages` to generate the metadata:

```
cd local-repo/Debian/bookworm

# download/copy .deb package files

dpkg-scanpackages . | gzip > Packages.gz
```

### RedHat/SLES repository metadata

First, install the `createrepo` package:

```
sudo yum install -y createrepo
```

Now you can use `createrepo` to generate the metadata:

```
cd local-repo/RedHat/8

# download/copy .rpm package files

createrepo .
```

## How TPA uses the local repository

### Copying the repository

TPA will use rsync to copy the contents of the repository directory, including the generated metadata, to a directory on target instances.

If rsync is not already available on an instance, TPA can install it (i.e., `apt-get install rsync` or `yum install rsync`). However, if you have set `use_local_repo_only`, the rsync package must be included in the local repo. If required, TPA will copy just the rsync package using scp and install it before copying the rest.

### Repository configuration

After copying the contents of the local repo to target instances, TPA will configure the destination directory as a local (i.e., path-based, rather than URL-based) repository.

If you provide, say, `example.deb` in the repository directory, running `apt-get install example` will suffice to install it, just like any package in any other repository.

### Package installation

TPA configures a repository with the contents that you provide, but if the same package is available from different repositories, it is up to the package manager to decide which one to install (usually the latest, unless you specify a particular version).

(However, if you set `use_local_repo_only: yes`, TPA will disable all other package repositories, so that instances can only use the packages that you provide in `local-repo`.)

### Installing from source

You can define a list of extensions to build and install from their Git repositories by setting `install_from_source` in `config.yml`:

```
cluster_vars:
  install_from_source:
    - name: ext
      git_repository_url: https://repo.example.com/ext.git
      git_repository_ref: dev/example

    - name: otherext
      git_repository_url: ssh://repo.example.com/otherext.git
      git_repository_ref: master
      source_directory: /opt/postgres/src/otherext
      build_directory: /opt/postgres/build/otherext
      build_commands:
        - "make -f /opt/postgres/src/otherext/Makefile install"
      build_environment:
        VAR: value
```

TPA will build and install extensions one by one in the order listed, so you can build extensions that depend on another (such as `pglogical` and `BDR`) by mentioning them in the correct order.

Each entry must specify a `name`, `git_repository_url`, and `git_repository_ref` (default: `master`) to build. You can use *Git credentials*

to authenticate to the Git repository; and also set `source_directory`, `build_directory`, `build_environment`, and `build_commands` as shown above.

Run `tpaexec deploy ... --skip-tags build-clean` in order to reuse the build directory when doing repeated deploys. (Otherwise the old build directory is emptied before starting the build.) You can also configure *Local source directories*

to speed up your development builds.

Whenever you run a source build, Postgres will be restarted.

## Build dependencies

If you're building from source, TPA will ensure that the basic Postgres build dependencies are installed. If you need any additional packages, mention them in *Installing packages*. For example

```
cluster_vars:
  packages:
    common:
      - golang-1.16
```

## Git credentials

This page explains how to clone Git repositories that require authentication.

This may be required when you change `postgres_git_url` to *Postgres source installation* or *Building and installing from source* to compile and install extensions.

You have two options to authenticate without writing the credentials to disk on the target instance:

- For an `ssh://` repository, you can add an SSH key to your local `ssh-agent`. Agent forwarding is enabled by default if you use `--install-from-source` (`forward_ssh_agent: yes` in `config.yml`).
- For an `https://` repository, you can export `TPA_GIT_CREDENTIALS=username:token` in your environment before running `tpaexec deploy`.

### Note

When deploying to Docker on macOS, you should use only `https://` repository URLs because Docker containers cannot be accessed by `ssh` from the host in this environment.

## SSH key authentication

If you are cloning an SSH repository and have an SSH keypair (`id_example` and `id_example.pub`), use SSH agent forwarding to authenticate on the target instances:

- **You need to run ``ssh-agent`` locally.** If your desktop environment does not already set this up for you (as most do—`pgrep ssh-agent` to check if it's running), run `ssh-agent bash` to temporarily start a new shell with the agent enabled, and run `tpaexec deploy` from that shell.
- **Add the required key(s) to the agent** with `ssh-add /path/to/id_example` (the private key file)
- **Enable SSH agent forwarding** by setting `forward_ssh_agent: yes` at the top level in `config.yml` before `tpaexec provision`. (This is done by default if you use `--install-from-source`.)

During deployment, any keys you add to your agent will be made available for authentication to remote servers through the forwarded agent connection.

Use SSH agent forwarding with caution, preferably with a disposable keypair generated specifically for this purpose. Users with the privileges to access the agent's Unix domain socket on the target server can co-opt the agent into impersonating you while authenticating to other servers.

## HTTPS username/password authentication

If you are cloning an HTTPS repository with a username and authentication token or password, just `export TPA_GIT_CREDENTIALS=username:token` in your environment before `tpaexec deploy`. During deployment, these credentials will be made available to any `git clone` or `git pull` tasks (only). They will not be written to disk on the target instances.

## Environment

You can set `target_environment` to specify environment variables that TPA should set on the target instances during deployment (e.g., to specify an HTTPS proxy, as shown below).

```
cluster_vars:
  target_environment:
    https_proxy: https://proxy.example:8080
```

TPA will ensure these settings are present in the environment (along with any others it needs) during deployment and the later execution of any cluster management commands.

These environment settings are not persistent, but you can instead use *The postgres Unix user* to set environment variables for the postgres user.

## Python environment

TPA decides which Python interpreter to use based on the *Distribution support* on a target instance. It will use Python 3 wherever possible, and fall back to Python 2 only when unavoidable.

The `tpaexec configure` command will set `preferred_python_version` according to the distribution.

| Distribution       | Python 2 | Python 3   |
|--------------------|----------|------------|
| Debian 12/bookworm | Yes      | Yes (3.11) |
| Debian 11/bullseye | Yes      | Yes (3.9)  |
| Debian 10/buster   | Yes      | Yes (3.7)  |
| Ubuntu 24.04/jammy | No       | Yes (3.12) |
| Ubuntu 22.04/jammy | No       | Yes (3.10) |
| Ubuntu 20.04/focal | No       | Yes (3.8)  |
| RHEL 9.x           | No       | Yes (3.9)  |
| RHEL 8.x           | No       | Yes (3.6)  |
| RHEL 7.x           | Yes      | No (3.6)   |

Ubuntu 20.04, 22.04 and RHEL 8.x can be used only with Python 3.

RHEL 7.x ships with Python 3.6, but the `librpm` bindings for system Python 3 are not available.

You can decide for other distributions whether you prefer `python2` or `python3`, but the default for new clusters is `python3`.

## Backwards compatibility

For compatibility with existing clusters, the default value of `preferred_python_version` is `python2`, but you can explicitly choose `python3` even on systems that were already deployed with `python2`.

```
cluster_vars:
  preferred_python_version: python3
```

TPA will ignore this setting on distributions where it cannot use Python 3.

## Configuring /etc/hosts

By default, TPA will add lines to /etc/hosts on the target instances with the IP address and hostname(s) of every instance in the cluster, so that they can use each other's names for communication within the cluster (e.g., in `primary_conninfo` for Postgres).

You can specify a list of `extra_etc_hosts_lines` too:

```
instances:
- Name: one
  ...
  vars:
    extra_etc_hosts_lines:
      - 192.0.2.1 acid.example.com
      - 192.0.2.2 water.example.com
```

If you don't want the default entries at all, you can specify the complete list of `etc_hosts_lines` for an instance instead, and only those lines will be added to /etc/hosts:

```
instances:
- Name: one
  ...
  vars:
    etc_hosts_lines:
      - 192.0.2.1 acid.example.com
      - 192.0.2.2 water.example.com
      - 192.0.2.3 base.example.com
```

If your /etc/hosts doesn't contain the default entries for instances in the cluster, you'll need to ensure the names can be resolved in some other way.

## Filesystem configuration

TPA allows you to define a list of `volumes` attached to each instance.

This list comprises both platform-specific settings that are used during provisioning and filesystem-level settings used during deployment.

First, `tpaexec provision` will use the information to create and attach volumes to the instance (if applicable; see platform-specific sections below for details). Then it will write a simplified list of volumes (containing only non-platform-specific settings) as a host var for the instance. Finally, `tpaexec deploy` will act on the simplified list to set up and mount filesystems, if required.

Here's a moderately complex example from an AWS cluster:

```
instances:
- Name: one
  ...
  volumes:
    - device_name: root
      volume_type: gp2
      volume_size: 32
    - device_name: /dev/xvdf
```

(continues on next page)

(continued from previous page)

```

volume_type: io2
volume_size: 64
iops: 5000
vars:
  volume_for: postgres_data
  encryption: luks
- device_name: /dev/xvdb
  ephemeral: ephemeral0
  vars:
    mountpoint: /mnt/scratch

```

In this example, the EC2 instance will end up with a 32GB EBS root volume, a 64GB io2 volume (provisioned-iops EBS volumes) mounted as `/opt/postgres/data`, and a `/tmp/scratch` filesystem provided by an instance-store (“ephemeral”) volume, whose number and size are determined by the instance type.

The details are documented in the section on AWS below, but settings like `volume_type` and `volume_size` are used during provisioning, while settings under `vars` like `volume_for` or `mountpoint` are written to the inventory for use during deployment.

## default\_volumes

Volumes are properties of an instance. You cannot set them in `cluster_vars`, because they contain platform-specific settings.

The *instance\_defaults*

mechanism makes special allowances for volume definitions. Since volume definitions in a large cluster may be quite repetitive (especially since we recommend that instances in a cluster be configured as close to each other as possible, you can specify `default_volumes` as shown here:

```

instance_defaults:
  default_volumes:
    - device_name: root
      volume_type: gp2
      volume_size: 32
    - device_name: /dev/xvdf
      volume_size: 100

instances:
- Name: one
  ...
- Name: two
  volumes:
    - device_name: /dev/xvdf
      volume_size: 64
    - device_name: /dev/xvdg
      volume_size: 64
  ...
- Name: three
  volumes:
    - device_name: /dev/xvdf

```

(continues on next page)

(continued from previous page)

```

    volume_type: none

- Name: four
  volumes: []

```

Here every instance will have a 32GB root volume and a 100GB additional volume by default (as is the case for instance `one`, which does not specify anything different). Instance `two` will have the same root volume, but it overrides `/dev/xvdf` to be 64GB instead, and has another 64GB volume in addition. Instance `three` will have the same root volume, but no additional volume because it sets `volume_type: none` for the default `/dev/xvdf`. Instance `four` will have no volumes at all.

An instance starts off with whatever is specified in `default_volumes`, and its `volumes` entries can override a default entry with the same `device_name`, remove a volume by setting `volume_type` to `none`, add new volumes with different names, or reject the defaults altogether.

(This behaviour of merging two lists is specific to `default_volumes`. If you set any other list in both `instance_defaults` and `instances`, the latter will override the former completely.)

## Platform AWS

On AWS EC2 instances, you can attach EBS volumes.

```

instances:

- Name: one
  ...
  volumes:
  - device_name: root
    volume_type: gp2
    volume_size: 32
    encrypted: yes
    ...
  - device_name: /dev/xvdf
    volume_type: io1
    volume_size: 32
    iops: 10000
    delete_on_termination: false
    ...
  - device_name: /dev/xvdg
    ephemeral: ephemeral0
    ...

```

TPA translates a `device_name` of `root` to `/dev/sda` or `/dev/xvda` based on the instance type, so that you don't need to remember (or change) which one to use.

The `volume_type` specifies the EBS volume type, e.g., `gp2` (for “general-purpose” EBS volumes), `io1` for provisioned-IOPS volumes (in which case you must also set `iops: 5000`), etc.

The `volume_size` specifies the size of the volume in gigabytes.

Set `encrypted: yes` to enable EBS encryption at rest. (This is an AWS feature, enabled by default in newly-generated TPA configurations, and is different from [LUKS encryption](#), explained below.)

Set `delete_on_termination` to `false` to prevent the volume from being destroyed when the attached instance is terminated (which is the default behaviour).

Set `ephemeral: ephemeralN` to use a physically-attached [instance store volume](#), formerly known as an ephemeral volume. The number, type, and size of available instance store volumes depends on the instance type. Not all instances have instance store volumes. Use instance store volumes only for testing or temporary data, and EBS volumes for any data that you care about.

For an EBS volume, you can also set `snapshot: snap-xxxxxxx` to attach a volume from an existing snapshot. Volumes restored from snapshots may be extraordinarily slow until enough data has been read from S3 and cached locally. (In particular, you can spin up a new instance with `PGDATA` from a snapshot, but expect it to take several hours before it is ready to handle your full load.)

If you set `attach_existing: yes` for a volume, and there is an existing unattached EBS volume with matching Name/type/size/iops, a new volume will not be created when launching the instance, but instead the existing one will be attached to the instance the first time it starts. Reattached EBS volumes do not suffer from the performance limitations of volumes created from snapshots.

## Platform bare

TPA has no control over what volumes may be attached to pre-provisioned `bare` instances, but if you define volumes with the appropriate `device_name`, it will handle `mkfs` and `mount` for the devices if required.

## Platform Docker

Docker containers can have attached volumes, but they are bind-mounted directories, not regular block devices. They do not need to be separately initialised or mounted. As such, the configuration looks quite different.

```
instances:
- Name: one
  platform: docker
  ...
  volumes:
  - /host/path/to/dir:/tmp/container/path:ro
  - named_volume:/mnt/somevol:rw
```

You may recognise these volume specifications as arguments to `docker run -v`.

The volumes are attached when the container is created, and there are no further actions during deployment.

## RAID arrays

On AWS EC2 instances, only RAID 0 is recommended by Amazon. You can create RAID volumes with a similar command:

```
sudo mdadm --create --verbose /dev/md0 --level=0 --name=MY_RAID --raid-devices=number_
↪of_volumes device_name1 device_name2
```

This example will attach the block device named `/dev/md0`. The handling of `volume_for` or `mountpoint` during deployment happens as the same as with any other volume. TPA will handle `mkfs` and `mount` for it.

```
- Name: one
...
volumes:
- device_name: /dev/md0
  vars:
    volume_for: postgres_data
```

## LUKS encryption

TPA can set up a LUKS-encrypted device:

```
instances:
- Name: one
  ...
  volumes:
  - device_name: /dev/xyz
    vars:
      encryption: luks
      luks_volume: mappedname
      volume_for: ...
```

If a volume with `encryption: luks` set is not already initialised, TPA will use `cryptsetup` to first `luksFormat` and then `luksOpen` it to map it under `/dev/mapper/mappedname` before handling filesystem creation as with any other device.

(To avoid any possibility of data loss, TPA will refuse to set up LUKS encryption on a device that contains a valid filesystem already.)

If you create a LUKS-encrypted `volume_for: postgres_data`, TPA will configure Postgres to not start automatically at boot. You can use `tpaexec start-postgres clustername` to mount the volume and start Postgres (and `stop-postgres` to stop Postgres and unmap the volume).

The LUKS passphrase is generated locally and stored in the vault.

## Filesystem creation and mounting

If any device does not contain a valid filesystem, it will be initialised with `mkfs`.

```
instances:
- Name: one
  ...
  volumes:
  - device_name: /dev/xyz
    vars:
      volume_for: ...
      fstype: ext4
      fsopts:
        - -cc
        - -m 2
      mountopts: defaults,relatime,nosuid
      readahead: 65536
      owner: root
      group: root
      mode: "0755"
```

You can specify the `fstype` (default: `ext4`), `fsopts` to be passed to `mkfs` (default: `none`), and `mountopts` to be passed to mount and written to `fstab` (see below).

TPA will set the `readahead` for the device to 16MB by default (and make the value persist across reboots), but you can specify a different value for the volume as shown above.

There are two ways to determine where a volume is mounted. You can either specify a `mountpoint` explicitly, or you

can set `volume_for` to `postgres_data`, `postgres_wal`, `postgres_tablespace` or `barman_data`, and TPA will translate the setting into an appropriate mountpoint for the system.

Once the mountpoint is determined, the device will be mounted there with the given `mountopts` (default: `defaults,noatime`). An entry will also be created for the filesystem in `/etc/fstab`.

You may optionally specify `owner`, `group`, or `mode` for the volume, and these attributes will be set on the mountpoint. Remember that at this very early stage of deployment, you cannot count on the `postgres` user to exist. In any case, TPA will (separately) ensure that any directories needed by Postgres have the right ownership and permissions, so you don't have to do it yourself.

## Uploading artifacts

You can define `artifacts` to create or copy files to target instances:

```
cluster_vars:
  artifacts:
    - type: path
      path: /some/target/path
      state: directory
      owner: root
      group: root
      mode: "0755"
    - type: file
      src: /host/path/to/file
      dest: /target/path/to/file
      owner: root
      group: root
      mode: "0644"
    - type: archive
      src: example.tar.gz
      dest: /some/target/path
    - type: directory
      src: /host/path/a/
      dest: /target/path/b/
```

The following types are supported:

- Use `path` to create or remove and change the ownership or mode of files and directories (takes the same parameters as Ansible's `file` module, which it uses internally)
- Use `file` to copy a file from the controller and set the ownership and mode (uses `copy`)
- Use `archive` to extract files from an archive to a specified location (uses `unarchive`)
- Use `directory` to rsync a directory from the controller to target instances (uses `synchronize`)

The example shows one entry for each of the above artifact types, but you can use these or any other parameters that the corresponding Ansible module accepts.

Copying files and directories to target instances is a common-enough need that this feature provides a convenient shortcut you can use instead of writing a *TPA hook*.

## ssh\_key\_file

By default, `tpaexec provision` will use `ssh-keygen` to generate a new SSH keypair for the cluster (into files named `id_cluster_name` and `id_cluster_name.pub` inside the cluster directory).

If you want to use an existing key instead, you can set `ssh_key_file` at the top level of `config.yml` to the location of an SSH private key file. The corresponding public key must be available with an extension of `.pub` at the same location:

```
ssh_key_file: ~/.ssh/id_rsa
```

(If this file does not already exist, it will be created by `ssh-keygen` during provisioning.)

## Managing SSH host keys

TPA generates a set of SSH host keys while provisioning a cluster. These keys are stored in the cluster directory, under the `hostkeys` subdirectory. These host keys are automatically installed into `/etc/ssh` on AWS EC2 instances and Docker containers.

By default, these host keys are not installed on *bare(-metal servers)*, but you can set `manage_ssh_hostkeys` to enable it:

```
instances:
- Name: one
  ...
  platform: bare
  vars:
    manage_ssh_hostkeys: yes
```

You must initially set up `known_hosts` in your cluster directory with correct entries, as described in the docs for *bare(-metal servers)*. TPA will replace the host keys during deployment.

The `manage_ssh_hostkeys` setting is meaningful only for bare instances. The generated host keys will be installed on all other instances.

## known\_hosts

TPA will add entries for every host and its public host keys to the global `ssh_known_hosts` file on every instance in the cluster, so that they can ssh to each other without host key verification prompts, regardless of whether they have `manage_ssh_hostkeys` set or not.

## Postgres source installation

TPA will compile and install Postgres from source if you set `postgres_installation_method` to `src`. This feature is meant for use in development and testing, and allows you to switch between packaged and source builds within an identically-configured cluster.

Even here, you do not need to change the defaults, which will give you a working cluster with debugging enabled.

## Git repository

The default settings will build and install Postgres from the community Git repository, using the `REL_xx_STABLE` branch corresponding to your `postgres_version`. You can specify a different repository or branch (any valid git reference) as follows:

```
cluster_vars:
  postgres_git_url: git://git.postgresql.org/git/postgresql.git
  postgres_git_ref: REL_12_STABLE
```

The default `git.postgresql.org` repository does not require authentication, but if necessary, you can use *Git credentials*

to authenticate to other repositories.

The repository will be cloned into `postgres_src_dir` (default: `/opt/postgres/src/postgres`), or updated with `git pull` if the directory already exists (e.g., if you are re-deploying).

## Build customisation

By default, TPA will configure and build Postgres with debugging enabled and sensible defaults in `postgres_build_dir` (default: `/opt/postgres/build/postgres`). You can change various settings to customise the build:

```
cluster_vars:
  postgres_extra_configure_env:
    CFLAGS: "-O3"
  postgres_extra_configure_opts:
    - --with-llvm
    - --disable-tap-tests
```

This will run `./configure` with the options in `postgres_extra_configure_opts` and the settings from `postgres_extra_configure_env` defined in the environment. Some options are specified by default (e.g., `--with-debug`), but can be negated by the corresponding `--disable-xxx` or `--without-xxx` options. Building `--without-openssl` is not supported.

If required, you can also change the following default build commands:

```
cluster_vars:
  postgres_make_command: "make -s"
  postgres_build_targets:
    - "all"
    - "-C contrib all"
  postgres_install_targets:
    - "install"
    - "-C contrib install"
```

Run `tpaexec deploy ... --skip-tags build-clean` in order to reuse the build directory when doing repeated deploys. (Otherwise the old build directory is emptied before starting the build.) You can also configure [Local source directories](#)

to speed up your development builds.

Whenever you run a source build, Postgres will be restarted.

## Additional components

Even if you install Postgres from packages, you can compile and install extensions from source. There's a separate page about how to configure [Building and installing from source](#).

If you install Postgres from source, however, you will need to install extensions from source as well, because the extension packages typically depend on the Postgres package(s) being installed.

## Package installation

There's a separate page about [Installing Postgres-related packages](#)

with `postgres_installation_method: pkg` (the default).

## Installing packages

TPA installs a batch of non-Postgres-related packages early during the deployment, then all Postgres-related packages together, and then packages for optional components separately. This page is about installing packages like `sysstat` or `strace`, which have no dependency on Postgres packages.

You can add entries to `packages` under `cluster_vars` or a particular instance's `vars` in `config.yml`:

```
cluster_vars:
  packages:
    common:
      - pkg1
      - pkg2
    Debian:
      - debpkg1
    RedHat:
      - rhpkg1
      - rhpkg2
    Ubuntu:
      - ubpkg1
    SLES:
      - slespkg1
```

In the example above, TPA will install its own list of `default_packages` and the packages listed under `packages.common` on every instance, and the remaining distribution-specific packages based on which distribution the instance is running. If any of these packages is not available, the deployment will fail.

Don't list any packages that depend on Postgres; use *Installing Postgres-related packages* instead.

## Optional packages

You can specify a list of `optional_packages` to install. They will be installed if they are available, and ignored otherwise. As with the other settings, the `common` entries apply to every instance, whereas any other lists apply only to instances running the relevant distribution.

```
optional_packages:
  common:
    - pkg1
    - pkg2
  Debian:
    - debpkg4
```

## Removing packages

You can specify a list of `unwanted_packages` that should be removed if they are installed.

```
unwanted_packages:
  common:
    - badpkg1
  Ubuntu:
    - badpkg2
```

## Running initdb

TPA will first create `postgres_data_dir` if it does not exist, and ensure it has the correct ownership, permissions, and SELinux context. Then, unless the directory already contains a `VERSION` file, it will run `initdb` to initialise `postgres_data_dir`.

You can use the *pre-initdb hook*

to execute tasks before `postgres_data_dir` is created and `initdb` is run. If the hook initialises `postgres_data_dir`, TPA will find the `VERSION` file and realise that it does not need to run `initdb` itself.

You can optionally set `postgres_initdb_opts` to a list of options to pass to `initdb`:

```
cluster_vars:
  postgres_locale: de_DE.UTF-8
  postgres_initdb_opts:
    - --data-checksums
```

We recommend always including the `--data-checksums` option (which is included by default).

TPA will set `TZ=UTC` in the environment, and set `LC_ALL` to the `postgres_locale` you specify, when running `initdb`.

## Separate configuration directory

By default, `postgres_conf_dir` is equal to `postgres_data_dir`, and the Postgres configuration files (`postgresql.conf`, `pg_ident.conf`, `pg_hba.conf`, and the include files in `conf.d`) are created within the data directory. If you change `postgres_conf_dir`, TPA will move the generated configuration files to the new location after running `initdb`.

## Installing Postgres-related packages

TPA installs a batch of non-Postgres-related packages early during the deployment, then all Postgres-related packages together, and then packages for optional components separately. This page is about installing packages like `pglogical` that depend on Postgres itself.

To install extra packages that depend on Postgres (e.g., `Postgis`), list them under `extra_postgres_packages` in `cluster_vars` or a particular instance's `vars` in `config.yml`:

```
cluster_vars:
  extra_postgres_packages:
    common:
      - postgres-pkg1
      - postgres-pkg2
    Debian:
      - postgres-deb-pkg1
    RedHat:
      - postgres11-rhpkg1
      - postgres11-rhpkg2
    Ubuntu:
      - ubpkg1
    SLES:
      - slespkg1
```

The packages listed under `packages.common` will be installed on every instance, together with the default list of Postgres packages, and any distribution-specific packages you specify.

There's a separate page about *Postgres source installation*.

## PG Backup API

If you set `enable_pg_backup_api: true` in `config.yml` or use the `--enable-pg-backup-api` command line option during configure, instances with the `barman` role will install `pg-backup-api` and set up an apache proxy for client cert authentication. This apache proxy will use an SSL CA generated for the cluster to generate its server and client certificates.

```
cluster_vars:
  enable_pg_backup_api: true
```

## PG Backup API package version

By default, TPA installs the latest available version of `pg-backup-api`.

The version of the `pg-backup-api` package that is installed can be specified by including `pg_backup_api_package_version: xxx` under the `cluster_vars` section of the `config.yml` file.

```
cluster_vars:
  ...
  pg_backup_api_package_version: 2.0.0-1.el8
  ...
```

You may use any version specifier that `apt` or `yum` would accept.

If your version does not match, try appending a `*` wildcard. This is often necessary when the package version has an epoch qualifier like `2:...`.

## Installation Options

`pg-backup-api` will be installed via packages by default, but you can also install from a git branch or a local directory. See [Building from source](#) and [Building and installing from source](#) for more details.

## Verify Setup

Run `pg-backup-api status` on the `barman` node running `pg-backup-api` - if you get “OK” back, the `pg-backup-api` service is running.

To test that the proxy is working, run

```
curl --cert /etc/tpa/pg-backup-api/pg-backup-user.crt \
  --key /etc/tpa/pg-backup-api/pg-backup-user.key \
  -X GET https://{hostname}/diagnose
```

If it's working, you'll get a large json output. You can compare this with the output of `barman diagnose`, they should match exactly.

## SSL Certificates

The root certificate will be copied to `/etc/tpa/pg-backup-api/` by default.

A client certificate and key (`pg-backup-user.crt` and `pg-backup-user.key`) will be generated for testing (through `tpaexec test`) or command line from the `barman` host. See [tpaexec test](#).

An apache proxy server certificate and key (`pg-backup-api.crt` and `pg-backup-api.key`) will also be generated

Each service needing to query the api will need to generate its own client certificate separately. PEM agent role, for instance, generates a client certificate during its setup when both `--enable-pem` and `--enable-pg-backup-api` (or `config.yml` equivalent) are used.

### Minor update for pg backup api using `tpaexec upgrade`

When trying to upgrade to a specific package version, ensure the `pg_backup_api_package_version` in `config.yml` is updated to reflect the desired version.

The desired version can also be passed as an extra argument to the `tpaexec upgrade` command with:

```
tpaexec upgrade <cluster_dir> -e pg_backup_api_package_version="<desired version>"
--components=pg-backup-api
```

Refer to the section on *package version selection and upgrade* for more information.

To select PG Backup API for upgrade, ensure the `--components` flag passed to the `tpaexec upgrade` command contains `pg-backup-api` (or `all` )

Refer to the section on *component selection for upgrade*

for more information.

### Setting `sysctl` values

By default, TPA sets various `sysctl` values on target instances, and includes them in `/etc/sysctl.conf` so that they persist across reboots.

You can optionally specify your own values in `sysctl_values`:

```
cluster_vars:
  sysctl_values:
    kernel.core_pattern: core.%e.%p.%t
    vm.dirty_bytes: 4294967296
    vm.zone_reclaim_mode: 0
```

Any values you specify will take precedence over TPA's default values for that variable (if any). The settings will first be added to `sysctl.conf` line-by-line, and finally loaded with `sysctl -p`.

Docker and `lxd` instances do not support setting `sysctls`, so TPA will skip this step altogether for those platforms.

## 7.5.5 Postgres configuration

### Creating Postgres databases

To create Postgres databases during deployment, add entries to the list of `postgres_databases` under `cluster_vars` or a particular instance's `vars` in `config.yml`:

```
cluster_vars:
  postgres_databases:
    - name: exampledb

    - name: complexdb
      owner: example
      encoding: UTF8
      lc_collate: de_DE.UTF-8
      lc_ctype: de_DE.UTF-8
      template: template0
      extensions:
        - name: hstore
        - name: dblink
      languages:
```

(continues on next page)

(continued from previous page)

```
- name: plperl
- name: plpython
tablespace: exampletablespace
```

The example above would create two databases (apart from any databases that TPA itself decides to create, such as `bdr_database`).

Each entry must specify the `name` of the database to create. All other attributes are optional.

The owner is `postgres` by default, but you can set it to any valid username (the users in *Creating Postgres users* will have been created by this time).

The `encoding`, `lc_collate`, and `lc_ctype` values default to the `postgres_locale` set at the time of running `initdb` (the default is to use the target system's `LC_ALL` or `LANG` setting). If you are creating a database with non-default locale settings, you will also need to specify `template: template0`.

You can optionally specify the default `tablespace` for a database; the tablespace must already exist (see *Creating Postgres tablespaces*).

You can specify optional lists of `extensions` and `languages` to create within each database (in addition to any extensions or languages inherited from the template database). Any packages required must be installed already, for example by including them in *Installing Postgres-related packages*.

TPA will not drop existing databases that are not mentioned in `postgres_databases`, and it may create additional databases if required (e.g., for BDR).

## Creating Postgres tablespaces

To create Postgres tablespaces during deployment, define their names and locations in `postgres_tablespaces` under `cluster_vars` or a particular instance's `vars` in `config.yml`.

If you *Filesystem configuration* with `volume_for: postgres_tablespace` set and a `tablespace_name` defined, they will be added as default entries to `postgres_tablespaces`.

```
cluster_vars:
  postgres_tablespaces:
    explicit:
      location: /some/path

instances:
- Name: example
  ...
  volumes:
  - device_name: /dev/xvdh
    ...
    vars:
      volume_for: postgres_tablespace
      tablespace_name: implicit
```

The example above would create two tablespaces: `explicit` (at `/some/path`) and `implicit` (at `/opt/postgres/tablespaces/implicit/tablespace_data` by default, unless you specify a different mountpoint for the volume).

Every `postgres_tablespace` volume must have `tablespace_name` defined; the tablespace location will be derived from the volume's mountpoint.

Every entry in `postgres_tablespaces` must specify a tablespace name (as the key) and its `location`. If you are specifying tablespace locations explicitly, do not put tablespaces inside PGDATA, and do not use any volume mountpoint directly as a tablespace location (`lost+found` will confuse some tools into thinking the directory is not empty).

By default, the tablespace owner is `postgres`, but you can set it to any valid username (the users in [Creating Postgres users](#)

will have been created by this time).

Streaming replicas must have the same `postgres_tablespace` volumes and `postgres_tablespaces` setting as their upstream instance

You can set the default tablespace for a database in [Creating Postgres databases](#).

## Configuring .pgpass

TPA creates `~postgres/.pgpass` by default with the passwords for the `postgres_user` (`postgres` or `enterprisedb` by default depending on the `postgres_flavour`) in it, for use between cluster instances.

You can set `pgpass_users` to create entries for a different list of users. Note that the `pgpass_users` list overrides default values, so the `postgres_user` (`postgres/enterprisedb`) is NOT included unless you explicitly include it in the `pgpass_users` list.

You can also include the `postgres/pgpass` role from hook scripts to create your own `.pgpass` file:

```
- include_role: name=postgres/pgpass
vars:
  pgpassfile: ~otheruser/.pgpass
  pgpass_owner: otheruser
  pgpass_group: somegroup
  pgpass_users:
    - xyzuser
    - pqruser
```

## The postgres Unix user

This page documents how the postgres user and its home directory are configured.

There's a separate page about how to create [Creating Postgres users](#).

## Shell configuration

TPA will install a `.bashrc` file and ensure that it's also included by the `.profile` or `.bash_profile` files.

It will set a prompt that includes the username and hostname and working directory, and ensure that `postgres_bin_dir` is in the `PATH`, and set `PGDATA` to the location of `postgres_data_dir`.

You can optionally specify `extra_bashrc_lines` to append arbitrary lines to `.bashrc`. (Use the YAML multi-line string syntax `>-` to avoid having to worry about quoting and escaping shell metacharacters.)

```
cluster_vars:
  extra_bashrc_lines:
    - alias la=ls\ -la
    - >-
      export PATH="$PATH":/some/other/dir
```

It will edit `sudoers` to allow `sudo systemctl start/stop/reload/restart/status postgres`, and also change `ulimits` to allow unlimited core dumps and raise the file descriptor limits.

## SSH keys

TPA will use `ssh-keygen` to generate and install an SSH keypair for the postgres user, and edit `.ssh/authorized_keys` so that the instances in the cluster can ssh to each other as postgres.

## TLS certificates

By default, TPA will generate a private key and a self-signed TLS certificate which are used by Postgres as the `ssl_key_file` and `ssl_cert_file` respectively. The files are named using the TPA cluster name (`cluster_name.key` and `cluster_name.crt`) and located in `/etc/tpa`. For more information, including how to provide your own key and certificate, see the documentation for [SSL configuration](#).

The size of self-signed TLS key can be modified adding the variable `postgres_rsa_key_size` to the `cluster_vars` section:

```
(...)
cluster_vars:
  postgres_rsa_key_size: 4096
```

By default, the self-signed TLS certificate has a period of 3650 days. The period can be adjusted by defining `openssl_certificate_validity` (in days) in the `cluster_vars` section. If the certificate needs to be replaced, use the `tpa_tls_renew_tls_certificates` variable on `tpaexec deploy`: `tpaexec deploy (...) -e tpa_tls_renew_tls_certificates=true`.

## Username

The `postgres_user` and `postgres_group` settings (both `postgres` by default) are used consistently everywhere. You can change them if you need to run Postgres as a different user for some reason.

## Creating Postgres users

To create Postgres users during deployment, add entries to the list of `postgres_users` under `cluster_vars` or a particular instance's `vars` in `config.yml`:

```
cluster_vars:
  postgres_users:
    - username: example

    - username: otheruser
      generate_password: true
      role_attrs:
        - superuser
        - replication
      granted_roles:
        - r1
        - r2
```

The example above would create two users (apart from any users that TPA itself decides to create, such as `repmgr` or `barman`).

Each entry must specify the `username` to create.

Any roles in the `granted_roles` list will be granted to the newly-created user.

The `role_attrs` list may contain certain [CREATE ROLE options](#)

such as `[NO] SUPERUSER`, `[NO] CREATEDB`, `[NO] LOGIN` (to create a user or a role) etc.

## Password generation

By default, TPA will generate a random password for the user, and store it in a vault-encrypted variable named `<username>_password` in the cluster's inventory. You can retrieve the value later:

```
tpaexec show-password ~/clusters/speedy example
```

Output:

```
beePh~iez6lie4thi5KaiG%eghaeT]ai
```

You cannot explicitly specify a password in `config.yml`, but you can store a different `<username>_password` in the inventory instead:

```
tpaexec store-password ~/clusters/speedy example --random
tpaexec show-password ~/clusters/speedy example
```

Output:

```
tpaexec store-password ~/clusters/speedy example
```

Enter password at the prompt.

Output:

```
Password:
```

```
tpaexec show-password ~/clusters/speedy example
```

Output:

```
terrible insecure password
```

If you don't want the user to have a password at all, you can set `generate_password: false`.

## 7.5.6 Commands

### **tpaexec archive-logs**

To create a log directory and archive logs from instances, run

```
tpaexec archive-logs <cluster-dir>
```

This will create a `logs/YYYYMMDDHHMMss/` directory in your cluster directory and download a `tar.gz` archive of all the files under `/var/log` on each instance in the cluster into a separate directory.

### **Prerequisites**

If you have an existing cluster you can run `tpaexec archive-logs` immediately. But if you are configuring a new cluster, you must at least *tpaexec provision* the cluster. You will get more logs if you also *tpaexec deploy* the cluster.

## Quickstart

```
tpaexec archive-logs ~/clusters/speedy
```

```
PLAY [Prepare local host archive] *****

TASK [Collect facts] *****
ok: [localhost]

TASK [Set time stamp] *****
ok: [localhost]

TASK [Create local log archive directory] *****
changed: [localhost]

PLAY [Archive log files from target instances] *****

...

TASK [Remove remote archives] *****
changed: [kinship]
changed: [khaki]
changed: [uncivil]
changed: [urchin]

PLAY RECAP *****
khaki                : ok=3    changed=3    unreachable=0    failed=0
kinship              : ok=3    changed=3    unreachable=0    failed=0
localhost            : ok=3    changed=1    unreachable=0    failed=0
uncivil              : ok=3    changed=3    unreachable=0    failed=0
urchin               : ok=3    changed=3    unreachable=0    failed=0
```

You can append `-v`, `-vv`, etc. to the command if you want more verbose output.

## Generated files

You can find the logs for each instance under the cluster directory:

```
~/clusters/speedy/logs/
`-- 220220306T185049
    |-- khaki-logs-20220306T185049.tar.gz
    |-- kinship-logs-20220306T185049.tar.gz
    |-- uncivil-logs-20220306T185049.tar.gz
    `-- urchin-logs-20220306T185049.tar.gz
```

Archive contents example:

```
khaki-logs
|-- anaconda
|   |-- anaconda.log
|   |-- dbus.log
|   |-- dnf.librepo.log
|   |-- hawkey.log
```

(continues on next page)

(continued from previous page)

```
| |-- journal.log
| |-- ks-script-ipdkisn0.log
| |-- ks-script-jr03uzns.log
| |-- ks-script-mh2iidvh.log
| |-- lvm.log
| |-- packaging.log
| |-- program.log
| |-- storage.log
| |-- syslog
| `-- X.log
|-- btmp
|-- dnf.librepo.log
|-- dnf.log
|-- dnf.rpm.log
|-- hawkey.log
|-- lastlog
|-- private
|-- tpaexec.log
`-- wtmp
```

### tpaexec download-packages

The purpose of the downloader is to provide the packages required to do a full installation of a TPA cluster from an existing configuration. This is useful when you want to ship packages to secure clusters that do not have internet access, or avoid downloading packages repeatedly for test clusters.

The downloader will download the full dependency tree of packages required, and the resulting package repository will include metadata files for the target distribution package manager, so can be used exclusively to build clusters. At this time package managers Apt and YUM are supported.

#### Note

The download-packages feature requires Docker to be installed on the TPA host. This is because the downloader operates by creating a container of the target operating system and uses that system's package manager to resolve dependencies and download all necessary packages. The required Docker setup for download-packages is the same as that for *Installing Docker*.

### Usage

An existing cluster configuration needs to exist which can be achieved using the `tpaexec configure` command. No specific options are required to use the downloader. See *Cluster configuration*

Execute the download-packages subcommand to start the download process. Provide the OS and OS version that should be used by the downloader.

```
tpaexec download-packages cluster-dir --os RedHat --os-version 8
```

This can also be expressed as a specific docker image. It is strongly recommended that you use one of the tpa images prefixed like the example below.

```
tpaexec download-packages cluster-dir --docker-image tpa/redhat:8
```

The downloader will place files downloaded in the directory `local-repo` by default. It is possible to download to alternative directory by using the option `--download-dir path`.

### Using the result

The contents of the `local-repo` directory is populated with a structure determined by ansible according to the OS contained in the docker image. For example, the docker image `tpa/redhat:8` would have the following:

```
cluster-dir/
`-- local-repo
    |-- RedHat
    |   |-- 8
    |       |-- *.rpm
    |       |-- repodata
    |           |-- *repodata-files*
```

You can use this in the cluster as is or copy it to a target control node. See *Managing clusters in a disconnected or air-gapped environment*. A *Creating and using a local repository* will be detected and used automatically by TPA.

### Cleaning up failed downloader container

If there is an error during the download process, the command will leave behind the downloader container running to help with debugging. For instance you may want to log in to the failed downloader container to inspect logs or networking. Downloader container is typically named `$cluster_name-downloader` unless it exceeds the allowed limit of 64 characters for the container name. You can check for the exact name by running `docker ps` to list the running containers and look for a container name that matches your cluster name. In most cases you can log in to the running container by executing `docker exec -it $cluster_name-downloader /bin/bash`. After the inspection, you can clean up the left over container by running the `download-packages` command with `--tags cleanup`. For example:

```
tpaexec download-packages cluster-dir --docker-image tpa/redhat:8 --tags cleanup
```

### TPA custom commands

You can define custom commands that perform tasks specific to your environment on the instances in a TPA cluster.

You can use this mechanism to automate any processes that apply to your cluster. These commands can be invoked against your cluster directory, like any built-in cluster management command. Having a uniform way to define and run such processes reduces the likelihood of errors caused by misunderstandings and operator error, or process documentation that was correct in the past, but has drifted away from reality since then.

Writing Ansible playbooks means that you can implement arbitrarily complex tasks; following the custom command conventions means you can take advantage of various facts that are set based on your `config.yml` and the cluster discovery tasks that TPA performs, and not have to think about details like connections, authentication, and other basic features.

This makes it much easier to write resilient, idempotent commands in a way that ad-hoc shell scripts (could be, but) usually aren't.

## Quickstart

- Create `commands/mycmd.yml` within your cluster directory
- Run `tpaexec mycmd /path/to/cluster`

## Example

Here's an example of a command that runs a single command on all instances in the cluster. Depending on the use-case, you can write commands that target different hosts (e.g., `hosts: role_postgres` to run only on Postgres instances), or run additional tasks and evaluate conditions to determine exactly what to do.

```
- --  
  
# Always start with this  
  
- import_playbook: "{{ tpa_dir }}/architectures/lib/init.yml"  
  tags: always  
  
- name: Perform custom command tasks  
  hosts: all  
  tasks:  
  - name: Display last five lines of syslog  
    command: tail -5 /var/log/syslog  
    become_user: root  
    become: yes
```

## TPA custom tests

You can easily define in-depth tests specific to your environment and application to augment TPA's *tpaexec test*.

We strongly recommend writing tests for any tasks, no matter how simple, that you would run on your cluster to reassure yourself that everything is working as you expect. Having a uniform and repeatable way to run such tests ensures that you don't miss out on anything important, whether you're dealing with a crisis or just doing routine cluster management.

If you write tests that target cluster instances by their configured role (or other properties), you can be sure that all applicable tests will be run on the right instances. No need to look up or remember how many replicas to check the replication status on, nor which servers are running pgbouncer, or any other such details that are an invitation to making mistakes when you are checking things by hand.

Tests must not make any significant changes to the cluster. If it's not something you would think of doing on a production server, it probably shouldn't be in a test.

## Quickstart

- Create `tests/mytest.yml` within your cluster directory
- Run `tpaexec test /path/to/cluster mytest`

You can also create tests in some other location and use them across clusters with the `--include-tests-from /other/path` option to `tpaexec test`.

(Run `tpaexec help test` for usage information.)

## Example

Here's how to write a test that is executed on all Postgres instances (note `hosts: role_postgres` instead of `hosts: all`).

You can use arbitrary Ansible tasks to collect information from the cluster and perform tests. Just write tasks that will fail if some expectation is not met (`assert`, `fail` ... `when`, etc.).

```
--
- name: Perform my custom tests
  hosts: role_postgres
  tasks:

    # Always start with this
    - include_role:
        name: test
        tasks_from: prereqs.yml

    # Make sure that the PGDATA/PG_VERSION file exists. (This is just a
    # simplified example, not something that actually needs testing.)
    - name: Perform simple test
      command: "test -f {{ postgres_data_dir }}/PG_VERSION"
      become_user: "{{ postgres_user }}"
      become: yes

    - name: Run pg_controldata
      command: >
        {{ postgres_bin_dir }}/pg_controldata {{ postgres_data_dir }}
      register: controldata
      become_user: "{{ postgres_user }}"
      become: yes

    # Write output to clusterdir/$timestamp/$hostname/pg_controldata.txt
    - name: Record pg_controldata output
      include_role:
        name: test
        tasks_from: output.yml
      vars:
        output_file: pg_controldata.txt
        content: |
          {{ controldata.stdout }}
```

You can use the builtin `output.yml` as shown above to record arbitrary test output in a timestamped test directory in your cluster directory.

Each test must be a complete Ansible playbook (i.e., a list of plays, not just a list of tasks). It will be imported and executed after the basic TPA setup tasks.

## Destructive tests

Tests should not, by default, make any significant changes to a cluster. (Even if they do something like creating a table to test replication, they must be careful to clean up after themselves.)

Any test that makes changes to a cluster that would be unacceptable on a production cluster **MUST** be marked as `destructive`. These may be tests that you run only in development, or during the initial cluster “burn in” process.

You can define “destructive” tests by setting `destructive: yes` when including `prereqs.yml` in your test:

```
- hosts: ...
  tasks:
  - include_role:
      name: test
      tasks_from: prereqs.yml
  vars:
    destructive: yes
```

If someone then runs `tpaexec test /path/to/cluster mytest`, they will get an error asking them to confirm execution using the `--destroy-this-cluster` option.

(Note: using `--destroy-this-cluster` signifies an awareness of the risk of running the command. It does not guarantee that the test will actually destroy the cluster.)

### Configuring the PGAI agent

!!!Note

TPA currently only supports using the PGAI Agent to integrate clusters with EDB Postgres AI Cloud Service. Using the agent to integrate with EDB Postgres AI Hybrid Manager, or for usage reporting, will be supported in a future release of TPA.

### PGAI agent package version

By default, TPA installs the latest available version of `beacon-agent`.

The version of the `beacon-agent` package that is installed can be specified by including `beacon_agent_package_version: xxx` under the `cluster_vars` section of the `config.yml` file.

```
cluster_vars:
  ...
  beacon_agent_package_version: 1.56.2-1
  ...
```

You may use any version specifier that `apt` or `yum` would accept.

If your version does not match, try appending a `*` wildcard. This is often necessary when the package version has an epoch qualifier like `2:...`.

### PGAI agent configuration

The PGAI agent configuration contains two parameters which must be set per-cluster, the access key and the project id.

The access key is kept encrypted in the cluster directory and can be set or read using `tpa's store-password` and `show-password` commands:

```
tpaexec store-password . beacon_agent_access_key
```

```
Password:
```

```
tpaexec show-password . beacon_agent_access_key
```

If the environment variable `BEACON_AGENT_ACCESS_KEY` is set when `tpaexec provision` is run, the access key is set from its value.

The project id is stored in `config.yml` under the `beacon_agent_project_id` key in `cluster_vars`. If the `--beacon_agent_project_id` argument is passed to `tpaexec configure` then its value is written to `config.yml` appropriately.

## Installing the PGAI agent

TPA installs the PGAI agent from EDB's repositories and creates an operating system user called `beacon` and a database user called `beacon`. A configuration file for the agent is written to `.beacon/beacon_agent.yml` in the `beacon` user's home directory.

## Running the PGAI agent

TPA installs a systemd service unit file to start the agent at boot-time, running as the `beacon` user.

## Compliance

TPA can generate configurations designed to make it easy for a cluster to comply with the STIG or CIS standards. If you pass `--compliance stig` or `--compliance cis` to `tpaexec configure`, TPA will:

- Check that other options are compatible with the appropriate standard.
- Add various entries to the generated `config.yml`, including marking that this is a cluster meant to comply with a particular standard and setting Postgres configuration as required by the standard.
- Adjust some deployment tasks to enforce compliance.
- Run checks at the end of deployment.

The deploy-time checks can be skipped by giving the option `--excluded_tasks=compliance` to `tpaexec deploy`. This feature is intended for testing only, when using a test system on which full compliance is impossible (for example, because SSL certificates are not available).

There are some situations in which TPA will intentionally fail to comply with the selected standard; these are documented under Exceptions below.

## STIG

STIG compliance is indicated by the `--compliance stig` option to `tpaexec configure`.

## Option compatibility

STIG compliance requires the `bare` platform and the `epas` flavour. It requires the RedHat OS with version 8 or 9.

## Settings in config.yml

The following entry is added to `cluster_vars` to use the SQL/Protect feature of EDB Postgres Advanced Server:

```
extra_postgres_extensions: [ sql_protect ]
```

The following entries are added to `cluster_vars` to force clients to use SSL authentication:

```
hba_force_hostssl: True
hba_force_certificate_auth: True
hba_cert_authentication_map: sslmap
```

The following entries are added to `cluster_vars` to set GUCs in `postgresql.conf`:

```
tcp_keepalives_idle: 10
tcp_keepalives_interval: 10
tcp_keepalives_count: 10
log_destination: "stderr"
postgres_log_file_mode: "0600"
```

The following entries are added to `postgres_conf_settings` in `cluster_vars` to set GUCs in `postgresql.conf`:

```
edb_audit: "xml"
edb_audit_statement: "all"
edb_audit_connect: "all"
edb_audit_disconnect: "all"
statement_timeout: 1000
client_min_messages: "ERROR"
```

## Deployment differences

During deployment, TPA will set connection limits for the database users it creates, corresponding to the number of connections that are needed for normal operation. As each user is set up, it will also check that an SSL client certificate has been provided for it.

## Providing client ssl certificates

STIG requires DOD-approved ssl certificates for client connections. These certificates can't be generated by TPA and therefore must be supplied. When setting up authentication for a user from a node in the cluster, TPA will look for a certificate/key pair on the node. The certificate and key should be in files called `<username>.cert` and `<username>.key` in the directory given by the `ssl_client_cert_dir` setting. The default for this setting is `/`, so the files would be, for example, `/barman.cert` and `/barman.key` when the `barman` user is being set up.

## Final checks

At the end of deployment, TPA will check that the server has FIPS enabled.

## Exceptions

If you select EFM as the failover manager, TPA will configure password authentication for the EFM user. This goes against the STIG requirement that all TCP connections use certificate authentication. The reason for this exception is that EFM does not support certificate authentication.

## CIS

CIS compliance is indicated by the `--compliance cis` option to `tpaexec configure`.

## Settings in config.yml

The following entries are added to `cluster_vars` to set GUCs in `postgresql.conf`:

```
log_connections: "on"
log_disconnections: "on"
```

The following entry is added to `cluster_vars` to enable required extensions:

```
extra_postgres_extensions: ["passwordcheck", "pgaudit"]
```

The following entry is added to `cluster_vars` to set the umask for the postgres OS user:

```
extra_bash_rc_lines: "umask 0077"
```

The following entries are added to `postgres_conf_settings` in `cluster_vars` to set GUCs in `postgresql.conf`:

```
log_error_verbosity: "verbose"
log_line_prefix: "%m [%p]: [%l-1] db=%d,user=%u,app=%a,client=%h "
log_replication_commands: "on"
temp_file_limit: "1GB"
```

## Final checks

At the end of deployment, TPA will check that the server has FIPS enabled.

## Exceptions

TPA does not support pgBackRest as mentioned in the CIS specification. Instead TPA installs Barman.

TPA does not install and configure `set_user` as required by the CIS specification. This is because preventing logon by the Postgres user would leave TPA unable to connect to, and configure, the database.

## etcd

`etcd` is a distributed, reliable key-value store commonly used for distributed coordination. TPA configures `etcd` primarily as the Distributed Configuration Store (DCS) when using Patroni as the failover manager for PostgreSQL clusters.

TPA automatically deploys and configures a 3-node `etcd` cluster by default when you enable Patroni in the M1 architecture.

## Installation

TPA installs the `etcd` package available from the configured system repositories. On SLES and RHEL-based systems, TPA automatically enables the PGDG `extras` repository to provide the `etcd` package.

## etcd package version

By default, TPA installs the latest available version of `etcd`.

The version of the `etcd` package that is installed can be specified by including `etcd_package_version: xxx` under the `cluster_vars` section of the `config.yml` file.

```
cluster_vars:
  ...
  etcd_package_version: 3.6.2*
  ...
```

You may use any version specifier that `apt` or `yum` would accept.

If your version does not match, try prepending a `*` wildcard. This is often necessary when the package version has an epoch qualifier like `2:...`.

## Configuration

TPA generates the `/etc/etcd/etcd.conf` file based on your `config.yml` settings and cluster topology. It configures peer and client communication URLs, data directory, and SSL certificates.

## Security

TPA provides robust options for securing communication with and within the `etcd` cluster. This is controlled primarily by the `etcd_ssl_enabled` and `etcd_authentication_mode` variables.

## TLS Encryption

Setting `etcd_ssl_enabled: true` enables TLS encryption for all `etcd` traffic. TPA will automatically:

- Generate a Cluster CA and server certificates (including appropriate IP and DNS Subject Alternative Names) for each `etcd` node.
- Configure `etcd` to use HTTPS for both peer-to-peer and client-server communication.
- Configure clients (`etcdctl`, Patroni) to connect via HTTPS and validate the server certificate using the CA.

## Client Authentication

Once TLS encryption is enabled (`etcd_ssl_enabled: true`), you can choose an authentication mode using the `etcd_authentication_mode` variable:

- **“none” (default):** No client authentication is performed. Communication relies solely on TLS encryption if `etcd_ssl_enabled` is `true`.
- **“basic” :** Requires clients to authenticate using a username and password. TPA automatically:
  - Creates an administrative `root` user with a generated password.
  - If using Patroni, creates a dedicated, least-privilege user for Patroni (configurable via `patroni_etcd_user`) with `readwrite` access only to its specific key prefix (e.g., `/tpa/cluster_name`).
  - Configures Patroni to use these dedicated credentials.
- **“mtls” :** Requires clients (including peer nodes) to present a valid TLS certificate signed by the trusted cluster CA. TPA automatically:
  - Configures `etcd` to require and validate client certificates for both peer (`ETCD_PEER_CLIENT_CERT_AUTH=true`) and client (`ETCD_CLIENT_CERT_AUTH=true`) connections.
  - Configures clients (`etcdctl`, Patroni) to present their client certificates for authentication.

!!!Note

TPA includes robust logic to handle transitions between different authentication modes during deployment or reconfiguration, ensuring the cluster remains stable. Configuration validation checks are performed early to prevent invalid combinations (e.g., `mtls` requires `etcd_ssl_enabled: true`). However, it’s currently not able to handle transition of values for `etcd_ssl_enabled` (i.e. from `false` to `true` or vice-versa).

You can set the following variables for `etcd`.

| Variable                    | Default value     | Description                                                                            |
|-----------------------------|-------------------|----------------------------------------------------------------------------------------|
| <code>etcd_peer_port</code> | <code>2380</code> | The TCP port <code>etcd</code> uses for peer-to-peer (server-to-server) communication. |

continues on next page

Table 10 – continued from previous page

| Variable                               | Default value         | Description                                                                                                                                                                                                                                                                          |
|----------------------------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>etcd_client_port</code>          | <code>2379</code>     | The TCP port <code>etcd</code> uses for client communication.                                                                                                                                                                                                                        |
| <code>etcd_ssl_enabled</code>          | <code>false*</code>   | Enable SSL/TLS encryption for all <code>etcd</code> communication. See Security. * <code>true</code> for new clusters via <code>tpaexec configure</code> .                                                                                                                           |
| <code>etcd_authentication_mode</code>  | <code>none*</code>    | Defines the client authentication mode ( <code>none</code> , <code>basic</code> , <code>mtls</code> ). Requires <code>etcd_ssl_enabled: true</code> for modes other than <code>none</code> . See Security. * <code>mtls</code> for new clusters via <code>tpaexec configure</code> . |
| <code>etcd_compaction_mode</code>      | <code>revision</code> | The automatic compaction mode ( <code>revision</code> or <code>periodic</code> ).                                                                                                                                                                                                    |
| <code>etcd_compaction_retention</code> | <code>10</code>       | The retention value for automatic compaction. For <code>revision</code> mode, this is the number of revisions to keep. For <code>periodic</code> mode, this is the time interval (e.g., <code>1h</code> ).                                                                           |

### Minor update for etcd using tpaexec upgrade

!!!Warning ETCDCTL\_API v2 is deprecated

etcd v3.4 introduced a breaking change, making ETCDCTL\_API v3 the default version, but allowed the v2 API to be enabled manually.

etcd v3.6 completely removed the `--enable-v2` flag, making it impossible to use the v2 API at all.

TPA upgrade ONLY supports upgrade of `etcd` versions using the v3 API.

Before using TPA to upgrade `etcd` in a cluster, ensure it has

migrated from the v2 API to the v3 API

The desired version can also be passed as an extra argument to the `tpaexec upgrade` command with:

```
tpaexec upgrade -e etcd_package_version="<desired version>" --components=etcd
```

Refer to the section on [package version selection and upgrade](#) for more information.

To select `etcd` for upgrade, ensure the `--components` flag passed to the `tpaexec upgrade` command contains `etcd` (or `all` )

Refer to the section on [component selection for upgrade](#)

for more information.

### Locale

For some platform images and environments it might be desirable to set the region and language settings.

By default, TPAexec will install the `en_US.UTF-8` locale system files. You can set the desired locale in your `config.yml` :

```
user_locale: en_GB.UTF-8
```

To find supported locales consult the output of the following command:

```
localectl list-locales
```

Or the contents of the file `/etc/locales.defs` on Debian or Ubuntu.

### Patroni cluster management commands

Patroni can be used as a single master failover manager with the M1 architecture using the following command options.

```
tpaexec configure cluster_name -a M1 --enable-patroni --postgresql 14
```

Or by setting the `config.yml` option

```
cluster_vars:
  failover_manager: patroni
```

TPA is able to deploy Patroni clusters using either `patroni` packages (from PGDG repositories) or `edb-patroni` packages (from EDB repositories). You can configure that through the `patroni_package_flavour` option under `cluster_vars` in the `config.yml`, which can also be set through the `--patroni-package-flavour` command-line argument. If no `patroni_package_flavour` is explicitly set, TPA will attempt to infer the flavour based on the configured repositories: if EDB repositories were configured, implicitly select `edb` flavour, otherwise implicitly select `community` flavour.

TPA `configure` will add 3 `etcd` nodes, and may add 2 `haproxy` nodes if you specify the option `--enable-haproxy`. `Etcd` is used for the Distributed Configuration Store (DCS). Patroni supports other DCS backends, but they are not currently supported by EDB or TPA.

As an alternative to HAProxy, you can use the `--enable-pgbouncer` option to configure PgBouncer in the Postgres nodes. PgBouncer will be configured to pool connections for the primary. Patroni will be configured to reconfigure PgBouncer upon failovers or switchovers in the cluster, so PgBouncer follows the new primary Postgres instance.

TPA uses Patroni's feature of converting an existing PostgreSQL cluster. This allows for TPA to initialise and manage configuration. Once the PostgreSQL cluster has been created, Patroni will take the management over. TPA will then remove any postgres configuration files used during setup.

Once set up, Postgres can continue to be managed using TPA and settings in `config.yml` for the cluster. You can also use Patroni interfaces, such as the command line `patronictl` and the REST API, but it is recommended to use TPA methods wherever possible.

### Patroni package version

By default, TPA installs the latest available version of Patroni.

The version of the Patroni package that is installed can be specified by including `patroni_package_version: xxx` under the `cluster_vars` section of the `config.yml` file.

```
cluster_vars:
  ...
  patroni_package_version: 4.0.0-1PGDG.rhel8`
  ...
```

You may use any version specifier that `apt` or `yum` would accept.

If your version does not match, try appending a `*` wildcard. This is often necessary when the package version has an epoch qualifier like `2:...`.

TPA provides robust options for securing the Patroni REST API. This is controlled primarily by the `patroni_ssl_enabled` and `patroni_authentication_mode` variables.

## TLS Encryption

Setting `patroni_ssl_enabled: true` enables TLS encryption for the REST API. TPA will automatically:

- Generate a TPA cluster CA and server certificates for each Patroni node.
- Configure the Patroni REST API to serve traffic over HTTPS.
- Configure clients (like `patronictl` and HAProxy) to connect via HTTPS and validate the server certificate using the CA.

## Client Authentication

Once TLS encryption is enabled (`patroni_ssl_enabled: true`), you can choose an authentication mode using the `patroni_authentication_mode` variable:

- **``basic`` (default):** Requires clients to provide the configured Patroni REST API username and password. TPA generates these credentials automatically.
- **``mtls`` :** Requires clients to present a valid TLS certificate signed by the trusted cluster CA. TPA automatically:
- Configures Patroni REST API to require and validate client certificates (`verify_client: required`).
- Configures clients (`patronictl`, HAProxy) with the necessary client certificates and keys.

## API Access Control (Allowlist)

For enhanced security, Patroni's REST API is automatically configured with an `allowlist`. This list restricts access to potentially harmful API endpoints (those using `POST`, `PUT`, `PATCH`, `DELETE`) to only the other Patroni nodes within the cluster, based on the hostnames defined in the TPA inventory. Read-only endpoints (like health checks) remain accessible from other sources (e.g., HAProxy).

These configuration variables can be used to control certain behaviours in the deployment of Patroni in TPA.

| Variable                                 | Default value                                          | Description                                                                                                                                                                                                                            |
|------------------------------------------|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>patroni_super_user</code>          | postgres                                               | User to create in postgres for superuser role.                                                                                                                                                                                         |
| <code>patroni_replication_user</code>    | replicator                                             | Username to create in postgres for replication role.                                                                                                                                                                                   |
| <code>patroni_restapi_user</code>        | patroni                                                | Username to configure for the patroni REST API.                                                                                                                                                                                        |
| <code>patroni_rewind_user</code>         | rewind                                                 | Username to create in postgres for <code>pg_rewind</code> function.                                                                                                                                                                    |
| <code>patroni_installation_method</code> | pkg                                                    | Install patroni from packages or source (e.g. git repo or local source directory if docker).                                                                                                                                           |
| <code>patroni_package_flavour</code>     | community if no EDB repository is configured, else edb | Whether to install <code>edb-patroni</code> package ( <code>edb</code> flavour, requires EDB repositories) or <code>patroni</code> package ( <code>community</code> flavour, requires PGDG and EPEL (RedHat based only) repositories). |
| <code>patroni_ssl_enabled</code>         | no                                                     | Whether to enable SSL for REST API and ctl connection. Will use the cluster SSL cert and CA if available.                                                                                                                              |
| <code>patroni_authentication_mode</code> | basic                                                  | Defines the client authentication mode ( <code>basic</code> or <code>mtls</code> ) for the REST API. Requires <code>patroni_ssl_enabled: true</code> .                                                                                 |

continues on next page

Table 11 – continued from previous page

| Variable                             | Default value | Description                                                                                                                                                                                                                                        |
|--------------------------------------|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>patroni_rewind_enabled</i>        | yes           | Whether to enable postgres rewind, creates a user defined by <i>patroni_rewind_user</i> and adds config section.                                                                                                                                   |
| <i>patroni_watchdog_enabled</i>      | no            | Whether to configure the kernel watchdog for additional split brain prevention.                                                                                                                                                                    |
| <i>patroni_dcs</i>                   | etcd          | What backend to use for the DCS. The only option is etcd at the moment.                                                                                                                                                                            |
| <i>patroni_dcs_namespace</i>         | /tpa          | The namespace within the DCS under which Patroni stores its cluster state.                                                                                                                                                                         |
| <i>patroni_listen_port</i>           | 8008          | REST API TCP port number                                                                                                                                                                                                                           |
| <i>patroni_etcd_user</i>             | etcd_patroni  | The username for the dedicated etcd user Patroni uses when <i>etcd_authentication_mode</i> is <i>basic</i> .                                                                                                                                       |
| <i>patroni_conf_settings</i>         | {}            | A structured data object with overrides for patroni configuration. Partial data can be provided and will be merged with the generated config. Be careful to not override values that are generated based on instance information known at runtime. |
| <i>patroni_dynamic_conf_settings</i> | {}            | Optional structured data just for DCS settings. This will be merged onto <i>patroni_conf_settings</i> .                                                                                                                                            |
| <i>patroni_repl_max_lag</i>          | None          | This is used in the haproxy backend health check only when <i>haproxy_read_only_load_balancer_enabled</i> is true. See <a href="#">Configuring lvm-direct for production</a> for possible values for <i>/replica?lag</i>                           |

## !!!Note

When deploying via Ansible Tower, changing *patroni\_etcd\_user* from its default value will cause the deployment to fail. This happens because TPA cannot handle dynamic secret names in its inventory. To avoid this issue, you must manually generate the password locally and push it to your configuration repository after provisioning and before deploying:

```
tpaexec store-password <your_custom_username> .
```

Configuration for patroni is built from three layers, starting with defaults set by the Patroni daemon, config loaded from the DCS, and finally from local configuration. The last can be controlled from either configuration file and overrides via the environment. TPA controls the configuration file and values are built up in this order.

DCS config to be sent to the API and stored in the bootstrap section of the config file:

- TPA vars for `postgres` are loaded into the DCS settings, see [postgresql.conf](#) . Some features are not supported, see notes below.
- Patroni defaults for DCS settings
- User supplied defaults in *patroni\_dynamic\_conf\_settings* , if you want to override any DCS settings you can do that here.

Local config stored in the YAML configuration file:

- `bootstrap.dcs` loaded from previous steps above.
- configuration enabled by feature flags, such as *patroni\_ssl\_enabled* , see table above.

- then finally overloaded from user supplied settings, the `patroni_conf_settings` option. If you want to change or add configuration not controlled by a feature flag then this is the best place to do it.

Please note that configuration is *merged* on top of configuration generated by TPA from cluster information, such as IP addresses, port numbers, cluster roles, etc. Exercise caution in what you override as this might affect the stable operation of the cluster.

As Patroni stores all postgres configuration in the DCS and controls how and when this is distributed to postgres, some features of TPA are incompatible with patroni:

- It is not possible to change the template used to generate `postgresql.conf` with the setting `postgres_conf_template`.
- You cannot change the location of Postgres config files with the setting `postgres_conf_dir`.

### Patroni configuration in TPA `config.yml`

You can override single values:

```
cluster_vars:
  patroni_conf_settings:
    bootstrap:
      dcs:
        ttl: 120
```

Or full blocks (with an example from Patroni docs):

```
cluster_vars:
  patroni_conf_settings:
    restapi:
      http_extra_headers:
        X-Frame-Options: SAMEORIGIN
        X-XSS-Protection: 1; mode=block
        X-Content-Type-Options: nosniff
      https_extra_headers:
        Strict-Transport-Security: max-age=31536000; includeSubDomains
```

If you want to negate a value or section that is present in the default TPA config vars you can set the value to `null`. This will cause patroni to ignore this section when loading the config file.

For example the default TPA config for `log` is

```
log:
  dir: /var/log/patroni
```

To turn off logging add this to `config.yml`:

```
cluster_vars:
  patroni_conf_settings:
    log: null
```

### Etcd Integration

When `etcd` is used as the DCS, TPA automatically configures Patroni to connect to it based on the `etcd` security configuration:

- **Protocol:** Uses `http` or `https` based on `etcd_ssl_enabled`.

- **CA Certificate:** Provides the cluster CA certificate path when TLS is enabled.
- **Authentication:**
  - If `etcd_authentication_mode` is `basic`, Patroni is configured with the username (`patroni_etcd_user`) and its generated password.
  - If `etcd_authentication_mode` is `mtls`, Patroni is configured with the paths to its client certificate and key.

TPA provides these minimal set of tools for managing Patroni clusters.

### Status

To see the current status of the TPA cluster according to Patroni run

```
tpaexec status cluster_name
```

### Switchover

To perform a switchover to a replica node (e.g. to perform maintenance) run the command

```
tpaexec switchover cluster_name new_primary
```

The `new_primary` argument must be the name of an existing cluster node that is currently running as a healthy replica. Checks will be performed to ensure this is true before a switchover is performed.

Once a switchover has been performed it is recommended that you run `deploy` and `test` to ensure a healthy cluster.

```
tpaexec deploy cluster_name
tpaexec test cluster_name
```

TPA will detect the current role of nodes during `deploy` regardless of what `config.yml` contains, for example if a different node is the leader.

### Minor update for patroni using `tpaexec upgrade`

When trying to upgrade to a specific package version, ensure the `patroni_package_version` in `config.yml` is updated to reflect the desired version.

The desired version can also be passed as an extra argument to the `tpaexec upgrade` command with:

```
tpaexec upgrade <cluster_dir> -e patroni_package_version="<desired version>" --
↪components=patroni
```

Refer to the section on [package version selection and upgrade](#) for more information.

To select Patroni for upgrade, ensure the `--components` flag passed to the `tpaexec upgrade` command contains `patroni` (or `all`)

Refer to the section on [component selection for upgrade](#)

for more information.

### pg\_hba.conf

The Postgres documentation explains the various options available in `pg_hba.conf`.

By default, TPA will generate a sensible `pg_hba.conf` for your cluster, to allow replication between instances, and connections from authenticated clients.

You can add entries to the default configuration by providing a list of `postgres_hba_settings` :

```
cluster_vars:
  postgres_hba_settings:
    - "# let authenticated users connect from anywhere"
    - hostssl all all 0.0.0.0/0 scram-sha-256
```

You can override the default `local all all peer` line in `pg_hba.conf` by setting `postgres_hba_local_auth_method: md5`.

If you don't want any of the default entries, you can change `postgres_hba_template` :

```
cluster_vars:
  postgres_hba_template: pg_hba.lines.j2
  postgres_hba_settings:
    - "# my lines of text"
    - "# and nothing but my lines"
    - "# ...not even any clients!"
    - hostssl all all 0.0.0.0/0 reject
```

You can even create `templates/my_hba.conf.j2` in your cluster directory and set:

```
cluster_vars:
  postgres_hba_template: my_hba.conf.j2
```

If you put any template files outside the cluster directory's `templates` subdirectory, make sure to specify the absolute path to the file:

```
# in the root of the cluster directory

cluster_vars:
  postgres_hba_template: "{{ cluster_dir }}/my_hba.conf.j2"
```

```
# in a subdirectory of the cluster directory that is NOT templates

cluster_vars:
  postgres_hba_template: "{{ cluster_dir }}/subdirectory/my_hba.conf.j2"
```

```
# in a directory outside of the cluster directory

cluster_vars:
  postgres_hba_template: /path/to/file/outside/cluster_dir/my_hba.conf.j2
```

If you just want to leave the existing `pg_hba.conf` alone, you can do that too:

```
cluster_vars:
  postgres_hba_template:
```

Although it is possible to configure `pg_hba.conf` to be different on different instances, we generally recommend a uniform configuration, so as to avoid problems with access and replication after any topology-changing events such as switchovers and failovers.

## pg\_ident.conf

You should not normally need to change `pg_ident.conf`, and by default, TPA will not modify it.

You can set `postgres_ident_template` to replace `pg_ident.conf` with whatever content you like.

```
cluster_vars:
  pg_ident_template: ident.j2
```

You will also need to create `templates/ident.j2` in the cluster directory:

```
{% for u in [unixuser1, unixuser2] %}
mymap {{ u }} dbusername
{% endfor %}
```

## pgd-cli

### pgd-cli package version

By default, TPA installs the latest available version of `pgd-cli`.

!!!Note Only `pgd-cli` v1 packages are supported for BDR-Always-ON architectures

TPA only installs `pgd-cli` in BDR-Always-ON clusters if `pgdcli` is explicitly added to the `role` list of a BDR node.

```
cluster_vars:
  pgdcli_package_version: 5.0.0-1
```

You may use any version specifier that `apt` or `yum` would accept.

If your version does not match, try prepending a `*` wildcard. This is often necessary when the package version has an epoch qualifier like `2:...`.

### Updating pgd-cli using `tpaexec upgrade`

!!!Note

The documentation refers to this component as `pgd-cli` but the name to pass to the `--components` list is `pgdcli`.

The desired version can also be passed as an extra argument to the `tpaexec upgrade` command with:

```
tpaexec upgrade <cluster_dir> -e pgdcli_package_version="<desired version>" \
  --components=pgdcli
```

Refer to the section on *package version selection and upgrade* for more information.

To select `pgd-cli` for upgrade, ensure the `--components` flag passed to the `tpaexec upgrade` command contains `pgdcli` (or `all`)

Refer to the section on *component selection for upgrade*

for more information.

## Adding Postgres extensions

### Default Postgres extensions

By default, TPA adds the following extensions to every Postgres database (and if needed, automatically adds the corresponding entries into shared preload libraries)

- `pg_stat_statements`
- `pg_freespacemap`
- `pg_visibility`
- `pageinspect`
- `pgstattuple`

## User defined extensions

Additional extensions can be configured within `config.yml`, by specifying the extension name, any required shared preload entries and the package containing the extension.

When adding extensions, be sure to include both the package name to `extra_postgres_packages` and the extension name under `extra_postgres_extensions` (or to the `extensions` list of a database defined under `postgres_databases`).

If the extension requires, add the shared preload entry name for the extension to the `preload_extensions` list. Note this name may differ from the extension name itself, so be sure to check the extension's own documentation.

Here is a quick example for an extension that requires to be added to the shared preload extension list with a different entry for extension and library name.

```
cluster_vars:
  [...]
  extra_postgres_packages:
    - postgresql-17-my-extension
  extra_postgres_extensions:
    - my-extension
  preload_extensions:
    - my_extension
```

The following sections provide further information.

- *Reconciling changes made outside of TPA*
- *Creating Postgres databases*
- *shared\_preload\_libraries*
- *Installing Postgres-related packages*

## TPA recognized extensions

The following list of extensions only require the extension name to be added in `config.yml` (either to `extra_postgres_extensions` OR to the `extensions` list of a database specified in `postgres_databases`) and TPA will automatically include the correct package and any required entries to `shared_preload_libraries`.

- `edb_pg_tuner`
- `query_advisor`
- `edb_wait_states`
- `sql_profiler`
- `pg_failover_slots`
- `sql_protect`
- `edb_stat_monitor`

- autocluster
- reldata
- bluefin
- postgis
- pgaudit
- passwordcheck

## postgresql.conf

TPA creates a `conf.d` directory with various `.conf` files under it, and uses `include_dir` in the main `postgresql.conf` to use these additional configuration files.

The Postgres configuration files (`postgresql.conf`, `pg_ident.conf`, and `pg_hba.conf`) and the included files under `conf.d` are always stored in `postgres_conf_dir`. This is the same as `postgres_data_dir` by default, but you can set it to a different location if you wish to keep the configuration separate from the data directory.

The main configuration mechanism is to set variables directly:

```
cluster_vars:
  temp_buffers: 16MB
  log_connections: on
  autovacuum_vacuum_cost_limit: -1
  effective_cache_size: 4GB
  max_connections: 300
  max_wal_senders: 32
```

TPA splits the configuration up into multiple files. The two main files are `0000-tpa.conf` and `0001-tpa_restart.conf`. These contain settings that require a server reload or restart to change, respectively. During deployment, TPA will write any changes to the correct file and reload or restart Postgres as required.

TPA may use other files in certain circumstances (e.g., to configure optional extensions), but you do not ordinarily need to care where exactly a given parameter is set.

You should never edit any of the files under `conf.d`, because the changes may be overwritten when you next run `tpaexec deploy`.

## postgres\_port

The Postgres port can be overwritten from its default value to another valid port number.

```
cluster_vars:
  postgres_port: 5433
```

This method is preferred over setting the port via `postgres_conf_settings`, since `postgres_port` ensures that the value is correctly shared and used across all components relying on this value to interact with the postgres cluster.

## postgres\_conf\_settings

TPA provides variables like `temp_buffers` and `maintenance_work_mem` that you can set directly for many, but not all, available `postgresql.conf` settings.

You can use `postgres_conf_settings` to set any parameters, whether recognised by TPA or not. You need to quote the value exactly as it would appear in `postgresql.conf`:

```
cluster_vars:
  effective_cache_size: 2GB
  postgres_conf_settings:
    effective_cache_size: 4GB
    authentication_timeout: 1min
    synchronous_standby_names: >-
      any 2 ("first", "second", "third")
    bdr.global_lock_statement_timeout: 60s
```

This is most useful with settings that TPA does not recognise natively, but you can use it for any parameter (e.g., `effective_cache_size` can be set as a variable, but `authentication_timeout` cannot).

These settings will be written to `conf.d/9900-role-settings.conf`, and therefore take priority over variables set in any other way.

If you make changes to values under `postgres_conf_settings`, TPA has no way to know whether the a reload is sufficient to effect the changes, or if a restart is required. Therefore it will always restart the server to activate the changes. This is why it's always preferable to use variables directly whenever possible.

### shared\_buffers

By default, TPA will set `shared_buffers` to 25% of the available memory (this is just a rule of thumb, not a recommendation). You can override this default by setting `shared_buffers_ratio: 0.35` to use a different proportion, or by setting `shared_buffers_mb: 796` to a specific number of MB, or by specifying an exact value directly, e.g., `shared_buffers: "2GB"`.

### effective\_cache\_size

By default, TPA will set `effective_cache_size` to 50% of the available memory. You can override this default by setting `effective_cache_size_ratio: 0.35` to use a different proportion, or by setting `effective_cache_size_mb: 796` to a specific number of MB, or by specifying an exact value directly, e.g., `effective_cache_size: "8GB"`.

### shared\_preload\_libraries

TPA maintains an internal list of extensions that require entries in `shared_preload_libraries` to work, and if you include any such extensions in `postgres_extensions`, it will automatically update `shared_preload_libraries` for you.

If you are using unrecognised extensions that require preloading, you can add them to `preload_extensions`:

```
cluster_vars:
  preload_extensions:
    - myext
    - otherext
```

Now if you add `myext` to `postgres_extensions`, `shared_preload_libraries` will include `myext`.

By default, `shared_preload_libraries` is set in `conf.d/8888-shared_preload_libraries.conf`.

Setting `shared_preload_libraries` directly as a variable is not supported. You should not normally need to set it, but if unavoidable, you can set a fully-quoted value under `postgres_conf_settings`. In this case, the value is set in `conf.d/9900-tpa_postgres_conf_settings.conf`.

## Postgres log

The default log file is defined as `/var/log/postgres/postgres.log`. If you need to change that, you can now set `postgres_log_file` in your `config.yml`:

```
cluster_vars:
  [...]
  postgres_log_file: /srv/fantastic_logs/pg_server.log
```

TPA will take care of creating the directories and rotate the log when needed.

## SSL configuration

By default, TPA will generate a private key and a self-signed TLS certificate which are used by Postgres as the `ssl_key_file` and `ssl_cert_file` respectively. The files are named using the TPA cluster name (`cluster_name.key` and `cluster_name.crt`) and located in `/etc/tpa`, resulting in the following default configuration in `0001-tpa_restart.conf`:

```
ssl_key_file=/etc/tpa/cluster_name.key
ssl_cert_file=/etc/tpa/cluster_name.crt
```

This is sufficient to ensure that traffic between clients and server is encrypted in transit.

To provide your own certificates, upload them to the target nodes as *Uploading artifacts*, then set the path by specifying the following cluster variables:

```
cluster_vars:
  ...
  artifacts:
    - type: file
      dest: /path/to/your_key.key
      src: /local/path/to/your_key.key
      owner: root
      group: root
      mode: "0644"
    - type: file
      dest: /path/to/your_cert.crt
      src: /local/path/to/your_cert.crt
      owner: root
      group: root
      mode: "0600"
  ssl_key_file: /path/to/your_key.key
  ssl_cert_file: /path/to/your_cert.crt
```

Alternatively, if you upload your key and certificate to the default location, TPA will use them instead of generating its own, and you do not need to specify `ssl_key_file` or `ssl_cert_file`. Note, however, that you must explicitly create `/etc/tpa` because it doesn't exist at the time artifacts are uploaded. The permissions and ownership of these files will be adjusted by TPA when the `postgres` user is created during deployment.

```
cluster_vars:
  ...
  artifacts:
    - type: path
      path: /etc/tpa
      state: directory
```

(continues on next page)

(continued from previous page)

```

owner: root
group: root
mode: "0755"
- type: file
  dest: /etc/tpa/cluster_name.key
  src: /local/path/to/your_key.key
  owner: root
  group: root
  mode: "0644"
- type: file
  dest: /etc/tpa/cluster_name.crt
  src: /local/path/to/your_cert.crt
  owner: root
  group: root
  mode: "0600"

```

**!!!Note Other SSL settings**

TPA does not specify `ssl_ca_file` or `ssl_crl_file` by default. To provide these files yourself you can do so using [Uploading artifacts](#) and by specifying the cluster variables of the same name.

There are two ways you can override anything in the TPA-generated configuration.

The first (and recommended) option is to use `ALTER SYSTEM`, which always takes precedence over anything in the configuration files:

```
ALTER SYSTEM SET bdr.global_lock_statement_timeout TO 60s;
```

You can also edit `conf.d/9999-override.conf`:

```
echo "bdr.global_lock_statement_timeout=60s" >> conf.d/9999-override.conf
```

All other files under `conf.d` are subject to be overwritten during deployment if the configuration changes, but TPA will never change `9999-override.conf` after initially creating the empty file.

Depending on which settings you change, you may need to execute `SELECT pg_reload_conf()` or restart the server for the changes to take effect.

**Generating postgresql.conf from scratch**

By default, TPA will leave the default (i.e., `initdb` -generated) `postgresql.conf` file alone other than adding the `include_dir`. You should not ordinarily need to override this behaviour, but you can set `postgres_conf_template` to do so:

```

cluster_vars:
  postgres_conf_template: pgconf.j2

```

Now the `templates/pgconf.j2` in your cluster directory will be used to generate `postgresql.conf`.

**tpaexec deprovision**

Deprovision destroys a cluster and associated resources.

For a cluster using the `aws` platform, it will remove the instances and all keypairs, policies, volumes, security groups, route tables, VPC subnets, internet gateways and VPCs which were set up for the cluster.

For a cluster using the `docker` platform, it will remove the containers, any `ccache` directories which were set up for source builds in the containers, and any `docker` networks which were set up for the cluster.

For all platforms, it will remove all the files created locally by `tpaexec provision`, including `ssh` keys, stored passwords, `ansible` inventory, and logs.

### tpaexec info

You can use the `info` command to output information about the TPA installation. Providing this information is valuable for troubleshooting.

### Usage

- Run `tpaexec info`

### Subcommands

- `tpaexec info version`

Displays current TPA version

- `tpaexec info platforms`

Displays available deployment platforms

- `tpaexec info architectures`

Displays available deployment architectures

- `tpaexec info platforms/<name>`

Displays information about a particular platform

- `tpaexec info architectures/<name>`

Displays information about a particular architecture

### Example Output

The `tpaexec info` command outputs the following:

```
# TPAexec 23.29

tpaexec=/opt/EDB/TPA/bin/tpaexec
TPA_DIR=/opt/EDB/TPA
PYTHON=/opt/EDB/TPA/tpa-venv/bin/python3 (v3.12.18, venv)
TPA_VENV=/opt/EDB/TPA/tpa-venv
ANSIBLE=/opt/EDB/TPA/tpa-venv/bin/ansible (v2.16.3)
Validated: ea844d1b90295597d080bbf824dbbc6954886cb54ffdb265c7c71b99bedee67b [OK]
```

### tpaexec reconfigure

The `tpaexec reconfigure` command reads `config.yml` and generates a revised version of it that changes the cluster in various ways according to its arguments.

## Arguments

As with other `tpaexec` commands, the cluster directory must always be given.

### Changing a cluster's architecture

The following arguments enable the cluster's architecture to be changed:

- `--architecture <architecture>` (required) The new architecture for the cluster. Accepts `PGD-Always-ON` and `PGD-X` as valid arguments
- `--pgd-proxy-routing <global|local>` (required for `PGD-Always-ON`) How PGD-Proxy is to route connections. See *PGD-Always-ON* for more information about the meaning of this argument.
- `--edb-repositories <repositories>` (optional) A space-separated list of EDB package repositories. It is usually unnecessary to specify this; `tpaexec configure` will choose a suitable repository based on the postgres flavour in use in the cluster.

After changing the architecture, run *Upgrading your cluster* to make the required changes to the cluster.

### Changing a cluster from PGD-Proxy to Connection Manager in PGD-Always-ON

- `--enable-connection-manager` (required) The option to allow migration from PGD-Proxy on a PGD 5.9+ cluster to the new builtin Connection Manager.

### Changing a cluster from 2q to EDB repositories

The `--replace-2q-repositories` argument removes any 2ndQuadrant repositories the cluster uses and adds EDB repositories as required to replace them.

After reconfiguring with this argument, run `[tpaexec deploy](tpaexec-deploy.md)` to make the required changes to the cluster.

## Output format

The following options control the form of the output:

- `--describe` Shows a description of what would be changed, without changing anything.
- `--check` Validates the changes that would be made and shows any errors or warnings that result from validation, without changing anything.
- `--output <filename>` Writes the output to a file other than `config.yml`.

## Sample invocation

```
tpaexec reconfigure ~/clusters/speedy\
  --architecture PGD-Always-ON\
  --pgd-proxy-routing local
```

## 7.6 Selective task execution

### 7.6.1 Using task selectors

You can tell TPA to run only a subset of the tasks that constitute a full deployment using the `--excluded_tasks` and `--included_tasks` options to `tpaexec deploy`. Each of these arguments is a string treated as a comma-separated list of selectors. Equivalently, you can set the `excluded_tasks` and `included_tasks` variables in `config.yml`,

either for the whole cluster or for the separate instances. In `config.yml`, you can use either a comma-separated string or a yaml list.

Tasks matched by `excluded_tasks` are always excluded. If you specify `included_tasks`, then non-matching tasks are implicitly excluded.

Some selectors may be used in either list, and some only in the `excluded_tasks` list, as detailed below. A separate set of selectors is available for `tpaexec test`.

## 7.6.2 Examples

To deploy without running barman-related tasks:

```
tpaexec deploy <clustername> --excluded_tasks=barman
```

To deploy running only repmgr-related tasks:

```
tpaexec deploy <clustername> --included_tasks=repmgr
```

To deploy without trying to set hostnames on the instances:

```
tpaexec deploy <clustername> --excluded_tasks=hostname
```

To prevent bootstrap and ssh tasks from ever running, put the following into `config.yml`:

```
cluster_vars:
  excluded_tasks:
    - bootstrap
    - ssh
```

## 7.6.3 Supported selectors for `tpaexec deploy`

The following selectors are supported for either inclusion or exclusion:

- barman

Tasks related to Barman.

- bdr

Tasks related to setting up BDR, including when it is as used within a PGD cluster. If this selector is excluded, TPA will still install and configure the extension as specified in `config.yml`, but won't create the node groups or try to join the nodes.

- create\_pgproxy\_system\_user

Tasks which creates the system user for pgd\_proxy

- create\_postgres\_system\_user

Tasks which creates the system user for postgres

- efm

Tasks related to EFM.

- etcd

Tasks related to etcd.

- first-backup

Tasks which ensure the minimum number of barman backups exist.

- haproxy

Tasks related to haproxy.

- harp

Tasks related to harp.

- patroni

Tasks related to patroni.

- pem-agent

Tasks related to the PEM agent.

- pem-server

Tasks related to the PEM server.

- pem-webserver

Tasks related to configuring the web server on a PEM server.

- pg-backup-api

Tasks related to Barman's Postgres backup API.

- pgbouncer

Tasks related to PgBouncer.

- pgd-proxy

Tasks related to PGD Proxy.

- pglogical

Tasks related to pglogical.

- pkg

Tasks which install packages using the system package manager.

- post-deploy

The post-deploy hook, if one is defined.

- postgres

Tasks related to postgres.

- replica

Tasks which are run and instances acting as postgres replicas.

- repmgr

Tasks related to repmgr.

- restart

Tasks which restart services

- sys

Tasks related to system setup before any tasks specific to postgres or related software.

- zabbix-agent

Tasks related to the zabbix agent.

The following selectors are supported only for exclusion:

- artifacts

Tasks related to *Uploading artifacts* .

- barman-clean

Tasks which clean up the Barman build directory if Barman is being built from source.

- barman-pre-config

The barman-pre-config hook, if one is defined.

- bdr-pre-node-creation

The bdr-pre-node-creation hook, if one is defined.

- bdr-post-group-creation

The bdr-post-group-creation hook, if one is defined.

- bdr-pre-group-join

The bdr-pre-group-join hook, if one is defined.

- bootstrap

Tasks which ensure that python and other minimal dependencies are present before the rest of the deploy runs. Exclude this only if you are sure you have manually installed the relevant requirements.

- build-clean

Tasks which clean up build directories for any software that is being built from source.

- build-configure

Tasks which configure any software that is being built from source.

- cloudinit

Tasks which are run only on hosts managed by cloud-init.

- commit-scopes

Tasks related to configuration of BDR commit scopes.

- config

Tasks which create config files.

- efm-pre-config

The efm-pre-config hook, if one is defined.

- efm-post-config

The efm-post-config hook, if one is defined.

- fs

Tasks related to setting up additional *Filesystem configuration* on instances.

- hostkeys

Tasks which set up *Managing SSH host keys* .

- hostname

Tasks which set the hostname.

- hosts

Tasks which *Configuring /etc/hosts*

- initdb

Tasks which run initdb.

- local-repo

Tasks which set up *Creating and using a local repository* .

- locale

Tasks which install *Locale* .

- openvpn

Tasks which set up OpenVPN.

- pg-backup-api-clean

Tasks which clean up the build directory if the Postgres backup API is being built from source.

- pgbouncer-config

Tasks which create configuration files for pgbouncer.

- pgpass

Tasks which create the *Configuring .pgpass* file.

- post-repo

The post-repo hook, if one is defined.

- postgres-clean

Tasks which clean up the build directory if postgres is being built from source.

- postgres-config

The postgres-config hook, if one is defined.

- postgres-config-final

The postgres-config-final hook, if one is defined.

- pre-deploy

The pre-deploy hook, if one is defined.

- pre-initdb

The pre-initdb hook, if one is defined.

- replication-sets

Tasks related to witness-only replication sets on a BDR3 cluster.

- repmgr-clean

Tasks which clean up the build directory if repmgr is being built from source.

- repmgr-configure

Tasks which configure repmgr if it is being built from source.

- repo

Tasks which set up package repositories.

- rsyslog

Tasks related to rsyslog.

- service

Tasks related to system services, including configuration and restarting.

- src

Tasks which build and install packages from source.

- ssh

Tasks related to setting up ssh between instances.

- systemctl

Tasks which set and reload systemctl settings.

- sysstat

Tasks related to the sysstat service.

- tpa

Tasks related to TPA's own files installed on instances.

- user

Tasks related to setting up system users.

- watchdog

Tasks related to the kernel watchdog on a patroni cluster.

## **7.6.4 Supported selectors for `tpaexec test`**

The following selectors apply only for execution of `tpaexec test` :

- camo

Tasks related to testing CAMO in a BDR or PGD cluster.

- ddl

Tasks related to testing DDL in a BDR or PGD cluster.

- fail

Tasks which abort tests if a problem is detected. Exclude this selector to run tests regardless of failures.

- pgbench

Tasks which run pgbench.

- sys

Tasks which run system-level tests.

- barman, bdr, haproxy, pg-backup-api, pgbouncer, pgd-proxy, postgres, repmgr,

Tasks which test the various software components.