



Trusted Postgres Architect

Release 23.22.0

EnterpriseDB UK Limited

Sep 12, 2023

CONTENTS

1	Introduction	3
1.1	What can TPA do?	3
1.2	How do I use it?	3
1.2.1	Configuration	4
1.2.2	Provisioning	4
1.2.3	Deployment	5
1.2.4	Testing	5
1.2.5	Incremental changes	5
1.2.6	Cluster management	5
1.2.7	Extensible through Ansible	5
1.3	It's just Postgres	6
1.4	Versioning in TPA	6
1.4.1	Backwards compatibility	6
1.5	Getting started	6
2	Trusted Postgres Architect release notes	7
2.1	Trusted Postgres Architect 23.22 release notes	7
2.2	Trusted Postgres Architect 23.21 release notes	8
2.3	Trusted Postgres Architect 23.20 release notes	8
2.4	Trusted Postgres Architect 23.19 release notes	9
2.5	Trusted Postgres Architect 23.18 release notes	9
2.6	Trusted Postgres Architect 23.17 release notes	10
2.7	Trusted Postgres Architect 23.16 release notes	10
2.8	Trusted Postgres Architect 23.15 release notes	11
2.9	Trusted Postgres Architect 23.14 release notes	11
2.10	Trusted Postgres Architect 23.13 release notes	11
2.11	Trusted Postgres Architect 23.12 release notes	12
2.12	Trusted Postgres Architect 23.1 to 23.11 release notes	13
2.12.1	TPA 23.11	13
2.12.2	TPA 23.10	13
2.12.3	TPA 23.9	13
2.12.4	TPA 23.8	14
2.12.5	TPA 23.7	14
2.12.6	TPA 23.6	15
2.12.7	TPA 23.5	15
2.12.8	TPA 23.4	15
2.12.9	TPA 23.3	15
2.12.10	TPA 23.2	16
2.12.11	TPA 23.1	17

3	TPA installation	19
3.1	Quickstart	19
3.1.1	Add repository and install TPA on Debian or Ubuntu	19
3.1.2	Add repository and install TPA on RHEL, Rocky, AlmaLinux or Oracle Linux	19
3.1.3	Install additional dependencies	20
3.1.4	Verify installation (run as a normal user)	20
3.2	Where to install TPA	20
3.3	Installing TPA packages	20
3.3.1	Add repository on Debian or Ubuntu	20
3.3.2	Add repository on RHEL, Rocky, AlmaLinux or Oracle Linux	20
3.3.3	Install on Debian or Ubuntu	21
3.3.4	Install on RHEL, Rocky, AlmaLinux or Oracle Linux	21
3.4	Setting up the TPA Python environment	21
3.5	Installing TPA without internet or network access (air-gapped)	21
3.5.1	Downloading TPA packages	22
3.5.2	Installing without access to a Python package index	22
3.6	Verifying your TPA installation	22
3.7	Upgrading TPA	22
3.8	Ansible community support	23
4	Creating a cluster	25
4.1	Cluster configuration	25
4.1.1	Quickstart	25
4.1.2	Configuration options	25
4.1.3	Software selection	27
4.1.4	Overrides	30
4.1.5	Ansible Tower	30
4.1.6	git repository	30
4.1.7	Examples	31
4.2	tpaexec provision	33
4.2.1	Prerequisites	34
4.2.2	Quickstart	34
4.2.3	Options	35
4.2.4	Accessing the instances	35
4.2.5	Generated files	36
4.3	tpaexec deploy	37
4.3.1	Prerequisites	37
4.3.2	Quickstart	37
4.3.3	Selective deployment	38
4.3.4	deploy.yml	38
4.4	tpaexec test	39
4.4.1	Quickstart	39
5	Architectures	41
5.1	PGD-Always-ON	41
5.1.1	Cluster configuration	41
5.2	BDR-Always-ON	43
5.2.1	Cluster configuration	43
5.3	M1	45
5.3.1	Default layout	45
5.3.2	Application and backup failover	45
5.3.3	Cluster configuration	46
6	aws	49

6.1	API access setup	49
6.2	Introduction	51
6.2.1	Networking	51
6.2.2	Instances	51
6.3	Configuration	52
6.3.1	Regions	52
6.3.2	VPC (required)	52
6.3.3	AMI (required)	53
6.3.4	Subnets (optional)	53
6.3.5	Security groups (optional)	54
6.3.6	Internet gateways (optional)	54
6.3.7	SSH keys (optional)	55
6.3.8	S3 bucket (optional)	55
6.3.9	Instance profile (optional)	55
7	bare(-metal servers)	57
7.1	SSH access	57
7.2	Distribution support	58
7.3	IP addresses	58
8	Docker	59
8.1	Synopsis	59
8.2	Operating system selection	59
8.3	Installing Docker	59
8.3.1	CgroupVersion	60
8.3.2	Permissions	60
8.3.3	Docker container privileges	60
8.3.4	Privileged containers	60
8.3.5	security_opts and the no-new-privileges flag	61
8.3.6	Linux capabilities flags	61
8.3.7	Docker storage configuration	61
8.4	Docker container management	62
9	Customizations	63
9.1	Cluster configuration	63
9.1.1	config.yml	63
9.1.2	Variables	64
9.1.3	Cluster variables	64
9.1.4	Instance variables	65
9.1.5	instance_defaults	66
9.1.6	Locations	66
9.2	Instance configuration	67
9.2.1	System-level configuration	67
9.2.2	Postgres	69
9.2.3	Other components	70
9.2.4	After starting Postgres	70
9.3	Building from source	71
9.3.1	Quickstart	71
9.3.2	Configuration	71
9.3.3	Local source directories	72
9.3.4	Rebuilding	73
9.4	TPA hooks	73
9.4.1	Summary	73
9.4.2	General-purpose hooks	74

9.4.3	PGD hooks	75
9.4.4	Other hooks	75
9.4.5	New hooks	76
10	Upgrading your cluster	77
10.1	Introduction	77
10.2	Configuration	77
10.3	Upgrading from BDR-Always-ON to PGD-Always-ON	78
10.4	PGD-Always-ON	79
10.5	BDR-Always-ON	79
10.6	M1	79
10.7	Package version selection	80
11	tpaexec switchover	81
11.1	Example	81
11.2	Architecture options	81
12	BDR/HAProxy server pool management	83
12.1	Example	83
13	tpaexec rehydrate	85
13.1	Prerequisites	85
13.2	Example	86
13.2.1	Change the configuration	86
13.2.2	Start the rehydration	86
13.3	Rehydrate in phases	87
13.4	Appendix	87
13.4.1	Using awscli to change volume attributes	87
14	TPA and Ansible Tower/Ansible Automation Controller	89
14.1	Preparing Tower for working with TPA	89
14.1.1	Setting up a TPA virtual environment	89
14.1.2	Creating the TPA_2Q_SUBSCRIPTION_TOKEN credential type	89
14.2	Setting up a cluster	90
15	TPA, Ansible, and sudo	91
15.1	Ansible sudo invocations	91
15.2	Recommendations	92
15.3	SSH and sudo passwords	92
15.4	sudo options	92
15.5	Logging	93
15.6	Local privileges	93
16	TPA - PuTTY Configuration guide	95
16.1	Key conversion	95
16.2	ConfigurePuTTY	96
17	Troubleshooting	97
17.1	Recreate python virtual environment	97
17.2	Strange AWS errors regarding credentials	97
17.3	Logging	98
17.4	Cluster test	98
17.5	TPA server test	98
17.6	Skipping or including specific tags	99

18	Reference	101
18.1	Running TPA in a Docker container	101
18.1.1	Quickstart	101
18.1.2	Installing Docker	102
18.2	Installing TPA from source	102
18.2.1	Quickstart	102
18.2.2	Step-by-step	103
18.2.3	Python 3.6+	103
18.2.4	Virtual environment options	104
18.3	Managing clusters in a disconnected or air-gapped environment	104
18.3.1	Preparation	104
18.3.2	Downloading packages	105
18.3.3	Copying packages to the target environment	105
18.3.4	Deploying in a disconnected environment	105
18.3.5	Updating in a disconnected environment	105
18.4	Distribution support	105
18.4.1	Debian x86	106
18.4.2	Ubuntu x86	106
18.4.3	RedHat x86	106
18.4.4	SLES	106
18.4.5	Platform-specific considerations	106
18.5	TPA capabilities and supported software	106
18.5.1	Supported software	107
18.6	EDB Postgres Distributed configuration	107
18.6.1	Introduction	107
18.6.2	Installation	107
18.6.3	Overview of cluster setup	108
18.6.4	Instance roles	108
18.6.5	Configuration settings	108
18.6.6	Miscellaneous notes	110
18.6.7	Older versions of PGD	110
18.7	Barman	110
18.7.1	Barman configuration	111
18.7.2	Backup scheduling	111
18.7.3	SSH keys	111
18.8	Configuring EFM	112
18.8.1	EFM configuration	112
18.8.2	efm_conf_settings	112
18.9	Configuring haproxy	113
18.9.1	Read-Only load-balancer	113
18.9.2	Server options	113
18.9.3	Example	114
18.10	Configuring HARP	114
18.10.1	Installing HARP	114
18.10.2	Configuring HARP	114
18.10.3	Consensus layer	116
18.10.4	Configuring a separate user for harp proxy	116
18.10.5	Custom SSL password command	116
18.10.6	Configuring the harp service	116
18.10.7	Configuring harp http(s) health probes	117
18.11	Configuring Postgres Enterprise Manager (PEM)	117
18.11.1	Supported architectures	118
18.11.2	PEM configuration	118
18.11.3	Shared PEM server	118

18.11.4	Connecting to the PEM UI	120
18.12	Configuring pgbouncer	120
18.12.1	Databases	120
18.12.2	Authentication	121
18.12.3	Example	121
18.13	Configuring pgd-proxy	121
18.13.1	Configuration	121
18.14	pglogical configuration	123
18.14.1	Introduction	124
18.14.2	Publications	125
18.14.3	Subscriptions	125
18.14.4	Configuration changes	126
18.14.5	Interaction with PGD	127
18.14.6	Limitations	127
18.15	Configuring repmgr	127
18.15.1	repmgr configuration	128
18.15.2	repmgr on PGD instances	128
18.16	How TPA uses 2ndQuadrant and EDB repositories	128
18.16.1	Package sources used by TPA	128
18.16.2	Authenticating with package sources	129
18.16.3	How sources are selected by default	129
18.16.4	Specifying EDB 2.0 repositories	129
18.16.5	Specifying 2ndQuadrant repositories	130
18.17	Configuring EDB Repos 2.0 repositories	130
18.18	Configuring 2ndQuadrant repositories	131
18.19	Configuring APT repositories	131
18.20	Configuring YUM repositories	132
18.21	Creating and using a local repository	133
18.21.1	Creating a local repository with TPA	133
18.21.2	Creating a local repository manually	133
18.21.3	Local repo layout	133
18.21.4	Populating the repository and generating metadata	134
18.21.5	How TPA uses the local repository	135
18.22	Installing from source	135
18.22.1	Build dependencies	136
18.23	Git credentials	136
18.23.1	SSH key authentication	137
18.23.2	HTTPS username/password authentication	137
18.24	Environment	137
18.25	Python environment	138
18.25.1	Backwards compatibility	138
18.26	Configuring /etc/hosts	138
18.27	Filesystem configuration	139
18.27.1	default_volumes	140
18.27.2	Platform AWS	141
18.27.3	Platform bare	142
18.27.4	Platform Docker	142
18.27.5	RAID arrays	142
18.27.6	LUKS encryption	143
18.27.7	Filesystem creation and mounting	143
18.28	Uploading artifacts	144
18.29	ssh_key_file	145
18.30	Managing SSH host keys	145
18.30.1	known_hosts	146

18.31	Postgres source installation	146
18.31.1	Git repository	146
18.31.2	Additional components	147
18.31.3	Package installation	147
18.32	Installing packages	147
18.32.1	Optional packages	148
18.32.2	Removing packages	148
18.33	Running initdb	148
18.33.1	Separate configuration directory	149
18.34	Installing Postgres-related packages	149
18.35	SSL Certificates	150
18.36	Setting sysctl values	150
18.37	Creating Postgres databases	151
18.38	Creating Postgres tablespaces	151
18.39	postgresql.conf	152
18.39.1	postgres_conf_settings	153
18.39.2	shared_buffers	153
18.39.3	effective_cache_size	153
18.39.4	shared_preload_libraries	154
18.39.5	Making changes by hand	154
18.39.6	Generating postgresql.conf from scratch	154
18.40	pg_hba.conf	155
18.41	pg_ident.conf	155
18.42	Configuring .pgpass	156
18.43	The postgres Unix user	156
18.43.1	Shell configuration	156
18.43.2	SSH keys	157
18.43.3	TLS certificates	157
18.43.4	Username	157
18.44	Creating Postgres users	157
18.44.1	Password generation	158
18.45	tpaexec archive-logs	158
18.45.1	Prerequisites	158
18.45.2	Quickstart	158
18.45.3	Generated files	159
18.46	tpaexec download-packages	160
18.46.1	Usage	160
18.46.2	Using the result	161
18.46.3	Cleaning up failed downloader container	161
18.47	TPA custom commands	161
18.47.1	Quickstart	162
18.47.2	Example	162
18.48	TPA custom tests	162
18.48.1	Quickstart	163
18.48.2	Example	163
18.48.3	Destructive tests	164
18.49	Shared instances	164
18.50	Locale	164
18.51	Patroni cluster management commands	165
18.51.1	Patroni configuration file settings	166
18.51.2	Status	168
18.51.3	Switchover	168
18.52	tpaexec info	168
18.52.1	Usage	168

18.52.2	Example Output	169
18.53	tpaexec reconfigure	169
18.53.1	Arguments	169
18.53.2	Sample invocation	169

© Copyright EnterpriseDB UK Limited 2015-2023 - All rights reserved.

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 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).

1.2 How do I use it?

To use TPA, you need to install it and run the `tpaexec setup` command. Follow the [installation instructions](#) for your platform.

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

- Generate a cluster *Cluster configuration*
- `:ref:`Provisioning <Provisioning>` servers (VMs, containers) to host the cluster`

- `:ref:Deployment <Deployment>` software to the provisioned instances
- `:ref:Testing <Testing>` the deployed cluster

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

[tpa]$ tpaexec configure clustername --architecture M1 --platform aws \
    --postgresql 14 \
    --failover-manager repmgr

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

[tpa]$ tpaexec provision clustername

# 3. Deployment: install and configure the necessary software

[tpa]$ tpaexec deploy clustername

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

[tpa]$ 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 *list of capabilities and supported software* .

1.2.1 Configuration

The *tpaexec configure*

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 *make any configuration changes to your cluster*, 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.

1.2.2 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).

1.2.3 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.

1.2.4 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.

1.2.5 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.

1.2.6 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.

1.2.7 Extensible through Ansible

TPA supports a *variety of configuration options*, 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

- `:ref:TPA custom commands <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.
- `:ref:TPA custom tests <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.

- `:ref:`Hook scripts <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.

1.3 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.

1.4 Versioning in TPA

TPA previously used a date-based versioning scheme whereby the major version was derived from the year. From version 23 we have moved to a derivative of semantic versioning. For historical reasons, we are not using the full three-part semantic version number. Instead TPA uses a two-part `major.minor` format. The minor version is incremented on every release, the major version is only incremented where required to comply with the backward compatibility principle below.

1.4.1 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.

1.5 Getting started

Follow the *TPA installation instructions* for your system, then *configure your first cluster* .

TRUSTED POSTGRES ARCHITECT RELEASE NOTES

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

Version	Release date
<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>tpaexec-configure</i>	•

2.1 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 GitHub .

2.2 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.

2.3 Trusted Postgres Architect 23.20 release notes

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 this OpenJDK bug .
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>pgdcli</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.

2.4 Trusted Postgres Architect 23.19 release notes

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 <i>generate_password: false</i> .
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> .

2.5 Trusted Postgres Architect 23.18 release notes

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 <i>tpaexec test</i> 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 6 – 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.

2.6 Trusted Postgres Architect 23.17 release notes

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 configure 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 configure 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 etcd changes, ensuring that they are now idempotent and avoiding unnecessary restarts of etcd on subsequent deployments.

2.7 Trusted Postgres Architect 23.16 release notes

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 pgd-proxy 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
Enhancement	Added new CLI abbreviations for Postgres flavour and version, for example <code>-postgresql 14</code> or <code>-edbpg 15</code>

continues on next page

Table 8 – continued from previous page

Type	Description
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

2.8 Trusted Postgres Architect 23.15 release notes

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.

2.9 Trusted Postgres Architect 23.14 release notes

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)

2.10 Trusted Postgres Architect 23.13 release notes

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> .

2.11 Trusted Postgres Architect 23.12 release notes

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>Preliminary support for configuring multi-region AWS clusters.</p><p>Multi-region clusters require manual setup of VPCs and VPC.</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>Ensure that nodes subscribe to bdr_child_group, if available.</p><p> 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.</p>
Bug fix	<p>Fail reliably with a useful error if Postgres doesn’t start. </p><p> 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. </p>
Bug fix	<p>Don’t call <i>bdr.alter_node_replication_sets()</i> on witnesses for BDR 4.3 and later.</p><p> This adjusts to a new restriction in BDR versions where witness nodes are not handled with a custom replication set configuration. </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> Sort endpoints in pgd-proxy config to avoid file rewrites. </p> <p> This will likely require a pgd-proxy restart on the next deploy (but it will avoid unnecessary future rewrites/restarts on subsequent deploys). </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.

2.12 Trusted Postgres Architect 23.1 to 23.11 release notes

2.12.1 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.

2.12.2 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

2.12.3 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.

2.12.4 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.

2.12.5 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.

2.12.6 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.

2.12.7 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.

2.12.8 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

2.12.9 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 dcs 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

2.12.10 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 Etc for the BDR-Always-ON Architecture When using 'harp_consensus_protocol:etcd', explicitly add 'etcd' to the role for each of the 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 ### 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
- 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

2.12.11 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 “`harpctl 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

Omitextraconnectionattributes(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).

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`

3.1 Quickstart

Login to [EDB Repos 2.0](#)

to obtain your token. Then execute the following command, substituting your token for `<your-token>` and replacing `<your-plan>` with one of the following according to which EDB plan you are subscribed: `enterprise`, `standard`, `community360`, `postgres_distributed`.

3.1.1 Add repository and install TPA on Debian or Ubuntu

```
curl -sLf https://downloads.enterprisedb.com/<your-token>/<your-plan>/setup.deb.sh |   
↪ sudo -E bash  
sudo apt-get install tpaexec
```

3.1.2 Add repository and install TPA on RHEL, Rocky, AlmaLinux or Oracle Linux

```
curl -sLf https://downloads.enterprisedb.com/<your-token>/<your-plan>/setup.rpm.sh |   
↪ sudo -E bash  
sudo yum install tpaexec
```

3.1.3 Install additional dependencies

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

3.1.4 Verify installation (run as a normal user)

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

More detailed explanations of each step are given below.

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 *deploying to Docker containers*

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.

3.3 Installing TPA packages

To install TPA, you must first subscribe to an EDB repository that provides it. The preferred source for repositories is EDB Repos 2.0.

Login to [EDB Repos 2.0](#)

to obtain your token. Then execute the following command, substituting your token for `<your-token>` and replacing `<your-plan>` with one of the following according to which EDB plan you are subscribed: `enterprise`, `standard`, `community360`, `postgres_distributed`.

3.3.1 Add repository on Debian or Ubuntu

```
curl -sLf https://downloads.enterprisedb.com/<your-token>/<your-plan>/setup.deb.sh |  
↪ sudo -E bash
```

3.3.2 Add repository on RHEL, Rocky, AlmaLinux or Oracle Linux

```
curl -sLf https://downloads.enterprisedb.com/<your-token>/<your-plan>/setup.rpm.sh |  
↪ sudo -E bash
```

Alternatively, you may obtain TPA from the legacy 2ndQuadrant repository. To do so, login to the EDB Customer Support Portal and subscribe to the [products/tpa/release](#)

by adding a subscription under Support/Software/Subscriptions, and following the instructions to enable the repository on your system.

Once you have enabled one of these repositories, you may install TPA as follows:

3.3.3 Install on Debian or Ubuntu

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

3.3.4 Install on RHEL, Rocky, AlmaLinux or Oracle Linux

```
sudo yum install tpaexec
```

This will install TPA into `/opt/EDB/TPA`. It will also ensure that other required packages (e.g., Python 3.6 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.

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
```

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

3.5.1 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.

If this is not possible, please contact EDB support and we will provide you with a download link or instructions appropriate to your subscription.

3.5.2 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.

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.

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.

3.8 Ansible community support

TPA now supports ansible community, you may choose to use it by using `--use-community-ansible` option during `tpaexec setup`, default will be to use the legacy 2ndQuadrant/ansible fork. This will change in a future release, support for 2ndQuadrant/ansible will be dropped and community ansible will become the new default.

notable difference:

- change the `--skip-flags` options to community behavior where a task will be skipped if part of the list given to the `--skip-tags` option even if it is also tagged with special tag `always`. TPA expects all tasks tagged with `always` to be run to ensure a complete deployment, therefor `--skip-tags` should not be used when using community ansible.

CREATING A CLUSTER

4.1 Cluster configuration

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

4.1.1 Quickstart

```
[tpa]$ 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 [configuration settings that affect the deployment](#).)

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

4.1.2 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`. For a complete list of architectures, run `tpaexec info architectures`.

Next, you must specify a flavour and version of Postgres 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 `--num-cascaded-replicas 3` to create a cluster with three cascaded replicas. 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.,

Strange AWS errors regarding credentials 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

By default, each cluster will be configured with a number of randomly selected `/28` subnets from the CIDR range `10.33.0.0/16`, depending on the selected architecture.

Specify `--network 192.168.0.0/16` to assign subnets from a different network.

Note: On AWS clusters, this corresponds to the VPC CIDR. See

Strange AWS errors regarding credentials documentation for details.

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

Specify `--no-shuffle-subnets` to allocate subnets from the start of the network CIDR range, without randomisation, e.g. `10.33.0.0/28`, then `10.33.0.16/28` and so on.

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.

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

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`.

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.

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.

4.1.3 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.

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

2ndQuadrant and EDB repositories

TPA can enable any 2ndQuadrant or EDB software repository that you have access to through a subscription.

By default, it will install the 2ndQuadrant public repository (which does not need a subscription) and add on any product repositories that the architecture may require (e.g., the PGD repository).

More detailed explanation of how TPA uses 2ndQuadrant and EDB repositories is available [Where to install TPA](#)

Specify `--2q-repositories source/name/maturity ...` or `--edb-repositories repository ...` to specify the complete list of 2ndQuadrant or EDB repositories to install on each instance in addition to the 2ndQuadrant public repository.

If any EDB repositories are specified, any 2ndQuadrant ones will be ignored.

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

To use these options, you must `export TPA_2Q_SUBSCRIPTION_TOKEN=xxx` or `export EDB_SUBSCRIPTION_TOKEN=xxx` before you run `tpaexec`. You can get a 2ndQuadrant token from the 2ndQuadrant Portal under “Company info” in the left menu, then “Company”. You can get an EDB token from 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 [Local repository support](#) has more details.

Software versions

Postgres flavour and version

TPA supports PostgreSQL, EDB Postgres Extended, and EDB Postgres Advanced Server (EPAS) versions 11 through 15.

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 [EDB repository subscription](#).

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.pgdg90+1`
2. `--repmgr-package-version 4.0.5-1.pgdg90+1`
3. `--barman-package-version 2.4-1.pgdg90+1`
4. `--pglogical-package-version '2.2.0*'`
5. `--bdr-package-version '3.0.2*'`
6. `--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:... .`

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 `tpa exec deploy` on nodes with RHEL8 and above (order derivative OSs which used `dnf`). For example, if `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. Use `2ndqpostgres` instead of `postgres` to build and install 2ndQPostgres. By default, this will build the appropriate `REL_nnn_STABLE` branch.

You may use `--install-from-source 2ndqpostgres pglogical3 bdr3` to build and install all three components from source, or just use `--install-from-source pglogical3 bdr3` to use packages for 2ndQPostgres, but build and install pglogical v3 and PGD v3 from source. By default, this will build the `master` branch of pglogical and 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.

4.1.4 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.

4.1.5 Ansible Tower

Use the `--use-ansible-tower` and `--tower-git-repository` options to create a cluster adapted for deployment with Ansible Tower. See [Ansible Tower](#) for details.

4.1.6 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.)

4.1.7 Examples

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

```
[tpa]$ tpaexec configure ~/clusters/speedy --architecture M1 \
--num-cascaded-replicas 2 --distribution Debian \
--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
[tpa]$
```

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

```
$ ls ~/clusters/speedy
total 8
drwxr-xr-x 2 ams ams 4096 Aug  4 16:23 commands
-rw-r--r-- 1 ams ams 1374 Aug  4 16:23 config.yml
lrwxrwxrwx 1 ams ams  51 Aug  4 16:23 deploy.yml ->
                                         /home/ams/work/2ndq/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: {}

cluster_rules:

-   cidr_ip: 0.0.0.0/0
    from_port: 22
    proto: tcp
    to_port: 22

-   cidr_ip: 10.33.76.176/28
    from_port: 0
    proto: tcp
    to_port: 65535

-   cidr_ip: 10.33.148.240/28
    from_port: 0
    proto: tcp
    to_port: 65535
ec2_ami:
  Name: debian-10-amd64-20210721-710
  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:
  enable_pg_backup_api: false
  failover_manager: repmgr
  postgres_flavour: postgresql
  postgres_version: 14
  preferred_python_version: python3
  use_volatile_subscriptions: false
```

locations:

- **Name:** main
az: us-east-1a
region: us-east-1
subnet: 10.33.76.176/28
- **Name:** dr
az: us-east-1b
region: us-east-1
subnet: 10.33.148.240/28

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:** upsets
backup: kayak
location: main
node: 1
role:
 - primary
- **Name:** zebra
location: main
node: 2
role:
 - replica

(continues on next page)

(continued from previous page)

```

upstream: upsets

- Name: kayak
  location: main
  node: 3
  role:
    - barman
    - log-server
    - monitoring-server
    - witness
  upstream: upsets
  volumes:
    - device_name: /dev/sdf
      encrypted: true
      vars:
        volume_for: barman_data
        volume_size: 128
        volume_type: gp2

- Name: queen
  location: dr
  node: 4
  role:
    - replica
  upstream: zebra

- Name: knock
  location: dr
  node: 5
  role:
    - replica
  upstream: zebra

```

The next step is to run *tpaexec provision*

or learn more about how to customise the configuration of *the cluster as a whole* or how to configure an *individual instance* .

4.2 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. *M1*) and platform (e.g. *Strange AWS errors regarding credentials*) 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).

4.2.1 Prerequisites

Before you can provision a cluster, you must generate the cluster configuration with *tpaexec configure* (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.

4.2.2 Quickstart

```
[tpa]$ tpaexec provision ~/clusters/speedy

PLAY [Provision cluster] *****
  ↳ *****
  ...

TASK [Set up EC2 instances] *****
  ↳ *****
changed: [localhost] => (item=us-east-1:quirk)
changed: [localhost] => (item=us-east-1:keeper)
changed: [localhost] => (item=us-east-1:zealot)
changed: [localhost] => (item=us-east-1:quaver)
changed: [localhost] => (item=us-east-1:quavery)
...

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

PLAY RECAP *****
  ↳ *****
localhost                : ok=128  changed=20  unreachable=0  failed=0

real    2m19.386s
user    0m51.819s
sys     0m27.852s
```

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.

4.2.3 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.

4.2.4 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.

```
[tpa]$ cd ~/clusters/speedy
[tpa]$ cat ssh_config
Host *
    Port 22
    IdentitiesOnly yes
    IdentityFile "id_speedy"
    UserKnownHostsFile "known_hosts tpa_known_hosts"
    ServerAliveInterval 60

Host quirk
    User admin
    HostName 54.227.207.189
Host keeper
    User admin
    HostName 34.229.111.196
Host zealot
    User admin
    HostName 18.207.108.211
Host quaver
    User admin
    HostName 54.236.36.251
Host quavery
    User admin
    HostName 34.200.214.150
[tpa]$ ssh -F ssh_config quirk
Linux quirk 4.9.0-6-amd64 #1 SMP Debian 4.9.82-1+deb9u3 (2018-03-02) x86_64

The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
Last login: Sat Aug 4 12:31:28 2018 from 136.243.148.74
admin@quirk:~$ sudo -i
root@quirk:~#
```

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.

4.2.5 Generated files

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

```
[tpa]$ 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 `ssh configfile`, 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`.

```
[tpa]$ cat inventory/00-speedy
[tag_Cluster_speedy]
quirk ansible_host=54.227.207.189 node=1 platform=aws
keeper ansible_host=34.229.111.196 node=2 platform=aws
zealot ansible_host=18.207.108.211 node=3 platform=aws
quaver ansible_host=54.236.36.251 node=4 platform=aws
quavery ansible_host=34.200.214.150 node=5 platform=aws

[tpa]$ cat inventory/group_vars/tag_Cluster_speedy/01-speedy.yml
cluster_name: speedy
cluster_tag: tag_Cluster_speedy
postgres_version: 15
tpa_version: v23.10-22-g30c1d5ea
tpa_2q_repositories: []
vpn_network: 192.168.33.0/24

[tpa]$ cat inventory/host_vars/zealot/02-topology.yml
role:

- barman

- log-server
```

(continues on next page)

(continued from previous page)

```

- openvpn-server
- monitoring-server
- witness
upstream: quirk

```

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.

4.3 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).

4.3.1 Prerequisites

Before you can run `tpaexec deploy`, you must have already run

`tpaexec configure` to generate the cluster configuration and then provisioned

the servers with *`tpaexec provision`*.

Before deployment, you must export `TPA_2Q_SUBSCRIPTION_TOKEN=xxx` to enable any 2ndQuadrant repositories that require subscription. You can use the subscription token that you used to [Add repository and install TPA on Debian or Ubuntu](#) itself. If you forget to do this, an error message will soon remind you.

4.3.2 Quickstart

```

[tpa]$ 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

```

(continues on next page)

(continued from previous page)

```

quavery          : ok=260  changed=88  unreachable=0  failed=0
quirk            : ok=262  changed=100  unreachable=0  failed=0

real    7m1.907s
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.

4.3.3 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:

```
[tpa]$ 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.)

4.3.4 `deploy.yml`

The deployment process is architecture-specific. Here’s an overview of the various *configuration settings that affect the deployment*. 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 *creating hooks*. This protects you from future implementation changes within a particular architecture.

4.4 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`.

4.4.1 Quickstart

```
[tpa]$ 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.

ARCHITECTURES

5.1 PGD-Always-ON

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

This architecture is valid for use with EDB Postgres Distributed 5 only and requires a subscription to *Configuring EDB Repos 2.0 repositories*.

5.1.1 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>tpaexec-configure</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-location</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.

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 *`tpaexec help configure-options`*.

5.2 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 a subscription to the legacy 2ndQuadrant repositories, and some options require a subscription to EDB Repos 1.0. See *`How TPA uses 2ndQuadrant and EDB repositories`*

for more detail on this topic.

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.

5.2.1 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>tpaexec-configure</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> .
<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 *tpaexec help configure-options* .

5.3 M1

A Postgres cluster with a primary and a streaming replica, one Barman server, and any number of additional replicas cascaded from the first one. 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 provide a token. See [How TPA uses 2ndQuadrant and EDB repositories](#)

for more detail on this topic.

5.3.1 Default layout

By default, the primary has one read-only replica attached in the same location; the replica, in turn, has one cascaded replica attached in a different location, where the Barman server is also configured to take backups from the primary.




Fig. 1: Cluster with cascading replication

If there is an even number of Postgres nodes, the Barman node is additionally configured as a witness. This ensures that the number of nodes is always odd, which is convenient when enabling automatic failover.

5.3.2 Application and backup failover

The M1 architecture implements failover management in that it ensures that a replica will be promoted to take the place of the primary should the primary become unavailable. However it *does not provide any automatic facility to reroute application traffic to the primary*. If you require, automatic failover of application traffic you will need to configure this at the application itself (for example using multi-host connections) or by using an appropriate proxy or load balancer and the facilities offered by your selected failover manager.

The above is also true of the connection between the backup node and the primary created by TPA. The backup will not be automatically adjusted to target the new primary in the event of failover, instead it will remain connected to the original primary. If you are performing a manual failover and wish to connect the backup to the new primary, you may simply re-run `tpaexec deploy`. If you wish to automatically change the backup source, you should implement this using your selected failover manager as noted above.

5.3.3 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
```

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>tpaexec-configure</i> .
One of: <code>* -failover-manager {efm, repmgr, patroni} * -enable-efm</code> <code>* -enable-repmgr *</code> <code>-enable-patroni</code>	Select the failover manager from <i>tpaexec-configure</i> , <i>tpaexec-configure</i> and <i>tpaexec-configure</i> .

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>-num-cascaded-replicas</code>	The number of cascaded replicas from the first replica.	Defaults to 1.
<code>-enable-haproxy</code>	2 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>-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> .

More detail about M1 configuration

You may optionally specify `--num-cascaded-replicas N` to request N cascaded replicas (including 0 for none; default: 1).

You may also specify any of the options described by *[tpaexec help configure-options](#)* .

- – title: Platforms —

TPA fully supports provisioning production clusters on AWS EC2.

6.1 API access setup

To use the AWS API, you must:

- `Obtain an access keypair <[https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_credentials\\_access-keys.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html)>`_
- `Add it to your configuration <[https://boto.readthedocs.org/en/latest/boto\\_config\\_tut.html](https://boto.readthedocs.org/en/latest/boto_config_tut.html)>`_

For example,

```
[tpa]$ cat > ~/.aws/credentials
[default]
aws_access_key_id = AKIAIOSFODNN7EXAMPLE
aws_secret_access_key = wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
```

The AMI 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
ec2>DeleteRouteTable
ec2>DeleteSecurityGroup
ec2>DeleteSubnet
ec2>DeleteVpc
ec2:DescribeImages
ec2:DescribeInstanceStatus
ec2:DescribeInstances
```

(continues on next page)

(continued from previous page)

```
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:RunInstances
ec2:TerminateInstances
iam:AddRoleToInstanceProfile
iam:CreateInstanceProfile
iam:CreateRole
iam>DeleteInstanceProfile
iam>DeleteRole
iam>DeleteRolePolicy
iam:GetInstanceProfile
iam:GetRole
iam:GetRolePolicy:
iam:ListGroup
iam:ListInstanceProfiles
iam:ListInstanceProfilesForRole
iam:ListRolePolicies
iam:ListRoles
iam:ListUsers
iam:PassRole
iam:PutRolePolicy
iam:RemoveRoleFromInstanceProfile
kms:CreateGrant
kms:GenerateDataKeyWithoutPlaintext
s3:GetObject
s3:ListAllMyBuckets
s3:ListBucket
s3:PutObject
s3:PutObjectAcl
```

6.2 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.

6.2.1 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).

6.2.2 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.

6.3 Configuration

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

6.3.1 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.

6.3.2 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:
```

(continues on next page)

(continued from previous page)

```

filters:
  vpc-id: vpc-nnn
us-east-2:
  Name: Example
  filters:
    [filter expressions]

```

6.3.3 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.)

This platform supports Debian 9 (stretch), RedHat Enterprise Linux 7, Rocky 8, Ubuntu 16.04 (Xenial), and SUSE Linux Enterprise Server 15.

6.3.4 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

```

6.3.5 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

- 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).

6.3.6 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
```

6.3.7 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
```

6.3.8 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.

6.3.9 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
```


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

7.1 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:

- `:ref:`Use a different ssh key <ssh_key_file>``
- `:ref:`Manage ssh host keys for bare instances <Managing SSH host keys>``

7.2 Distribution support

TPA will try to detect the distribution running on target instances, and fail if it is not supported. TPA currently supports Debian (8/9/10; or jessie/stretch/buster), 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.

7.3 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.)

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.

8.1 Synopsis

Just select the platform at configure-time:

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

8.2 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.

TPA does not support Debian 8 (jessie) or Ubuntu 16.04 (xenial) for Docker containers, because of bugs in the old version of systemd shipped on those distributions.

8.3 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 for Mac”](#) and launch Docker from the application menu.

8.3.1 CgroupVersion

Support for CgroupVersion 2 is not fully baked yet for docker sdk in ansible and related tooling. So while we recommend using a recent version of docker; we rely on CgroupVersion 1 until version 2 is fully supported. Instructions below suggest the changes to switch to CgroupVersion 1 if your platform uses CgroupVersion 2 by default.

On Linux:

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

8.3.2 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:

```
[tpa]$ 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.

8.3.3 Docker container privileges

8.3.4 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.

See `man capabilities` for details on Linux capabilities flags.

8.3.5 security_opts and the no-new-privileges flag

tpaexec can start docker containers in a restricted mode where processes cannot increase their privileges. `se-tuid` 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.

8.3.6 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
```

8.3.7 Docker storage configuration

Caution : 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 <https://docs.docker.com/storage/storagedriver/>`_`
- ``Configuring lvm-direct for production <https://docs.docker.com/storage/storagedriver/device-mapper-driver/#configure-direct-lvm-mode-for-production>`_`

8.4 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:

```
[tpa]$ tpaexec start-containers clustername  
[tpa]$ 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:

```
[tpa]$ tpaexec list-containers clustername
```

CUSTOMIZATIONS

9.1 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 `tpaexec configure` 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

variables you can set to customise the deployment process .

9.1.1 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

- node: 2
  Name: two
  role: replica
  platform: docker
```

(continues on next page)

(continued from previous page)

```

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 *tpaexec* configure

options translate to `config.yml` variables in some way. A single option may affect several variables (e.g., `--bdr-version` could set `postgres_version`, `tpa_2q_repositories`, `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.

9.1.2 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.

9.1.3 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
  tpa_2q_repositories:
    - products/bdr3/release
    - products/pglogical3/release
  postgres_conf_settings:
    bdr.trace_replay: true

```

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`.

9.1.4 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`.

9.1.5 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`.

9.1.6 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.

9.2 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 [an overview of configuring a cluster](#) 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.

9.2.1 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 configuration below.

- `:ref:Distribution support <Distribution support>`
- `:ref:Setting up the TPA Python environment <Setting up the TPA Python environment>`
`↪`(`preferred_python_version`)`
- `:ref:Environment variables <Environment>` (e.g., ``https_proxy``)

Package repositories

You can use the *pre-deploy hook*

to execute tasks before any package repositories are configured.

- `:ref:Configure YUM repositories <Configuring YUM repositories>`

(for RHEL, Rocky and AlmaLinux)

- :ref:`Configure APT repositories <Configuring APT repositories>`

(for Debian and Ubuntu)

- :ref:`Configure 2ndQuadrant and EDB repositories <How TPA uses 2ndQuadrant and EDB repositories>`

(on any system)

- :ref:`Configure a local package repository <Creating and using a local repository>`

(to ship packages to target instances)

You can use the *post-repo* hook

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:

- :ref:`Install non-Postgres 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

- :ref:`Create and mount filesystems <Filesystem configuration>` (including RAID,  LUKS setup)
- :ref:`Upload artifacts <Uploading artifacts>` (files, directories, tar archives)
- :ref:`Set sysctl values <Setting sysctl values>`
- :ref:`Configure /etc/hosts <Configuring /etc/hosts>`
- :ref:`Manage ssh\_known\_hosts entries <Managing SSH host keys>`
- :ref:`Add system locale <Locale>`

9.2.2 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 *configure options* 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:

- `:ref:`Install Postgres and Postgres-related packages <Installing Postgres-related_↵packages>``

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

- `:ref:`Build and install Postgres and extensions from source <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

- `:ref:`Configure the postgres Unix user <The postgres Unix user>``
- `:ref:`Run initdb to create the PGDATA directory <Running initdb>``
- `:ref:`Configure pg_hba.conf <pg_hba.conf>``
- `:ref:`Configure pg_ident.conf <pg_ident.conf>``
- `:ref:`Configure postgresql.conf <postgresql.conf>``

You can use the *postgres-config hook*

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.

9.2.3 Other components

- [:ref:`Configure Barman <Barman>`](#)
- [:ref:`Configure pgbouncer <Configuring pgbouncer>`](#)
- [:ref:`Configure haproxy <Configuring haproxy>`](#)
- [:ref:`Configure HARP <Configuring HARP>`](#)
- [:ref:`Configure EFM <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.

9.2.4 After starting Postgres

- [:ref:`Create Postgres users <The postgres Unix user>`](#)
- [:ref:`Create Postgres tablespaces <Creating Postgres tablespaces>`](#)
- [:ref:`Create Postgres databases <Creating Postgres databases>`](#)
- [:ref:`Configure pglogical replication <pglogical configuration>`](#)
- [:ref:`Configure .pgpass <Configuring .pgpass>`](#)

You can use the *postgres-config-final hook*

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

- [:ref:`Configure BDR <EDB Postgres Distributed configuration>`](#)

You can use the *post-deploy hook*

to execute tasks after the deployment process has completed.

9.3 Building from source

TPA can build Postgres and other required components from source and deploy a cluster with exactly the same configuration as with the default packaged installation. This makes it possible to deploy repeatedly from source to quickly test changes in a realistic, fully-configured cluster that reproduces every aspect of a particular setup, regardless of architecture or platform.

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.

9.3.1 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.

9.3.2 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:

- `:ref: `Building Postgres from source` <Postgres source installation>``
- `:ref: `Building extensions with `install_from_source` <Installing from source>``

9.3.3 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.)

9.3.4 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.

9.4 TPA hooks

TPA can set up fully-functional clusters with no user intervention, and already provides a broad variety of *settings to control your cluster 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.

9.4.1 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.

9.4.2 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 *apt_repositories* or *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
run `initdb` to create `PGDATA` 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.

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.

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).

9.4.3 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.

9.4.4 Other hooks

postgres-pre-update, postgres-post-update

The *upgrade* 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).

9.4.5 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.

- – title: Cluster management —

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.)

10.1 Introduction

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.

This command will try to perform the upgrade with minimal disruption to cluster operations. The exact details of the specialised upgrade process depend on the architecture of the cluster, as documented below.

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.

10.2 Configuration

In many cases, minor-version upgrades do not need changes to `config.yml`. Just run `tpaexec upgrade`, and it will upgrade to the latest available versions of the installed packages in a graceful way (what exactly that means depends on the details of the cluster).

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.

10.3 Upgrading from BDR-Always-ON to PGD-Always-ON

To upgrade from BDR-Always-ON to PGD-Always-ON (that is, from BDR4 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 [the command's own documentation](#) . 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.

10.4 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.
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 pgd-proxy doesn't send any connections to it.
 - Stops, updates, and restarts postgres.
 - Unfences the node so it can receive connections again.
 - Updates pgbouncer, pgd-proxy, and pgd-cli, as applicable for this node.

10.5 BDR-Always-ON

For BDR-Always-ON clusters, the upgrade process goes through the cluster instances one by one and does the following:

1. Tell haproxy the server is under maintenance.
2. If the instance was the active server, request pgbouncer to reconnect and wait for active sessions to be closed.
3. Stop Postgres, update packages, and restart Postgres.
4. Finally, mark the server as “ready” again to receive requests through haproxy.

PGD logical standby or physical replica instances are updated without any haproxy or pgbouncer interaction. Non-Postgres instances in the cluster are left alone.

You can control the order in which the cluster's instances are updated 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 update by listing the active PGD primary instances last, so that the shadow servers are updated first.

If your environment requires additional actions, the *postgres-pre-update* and *postgres-post-update hooks* allow you to execute custom Ansible tasks before and after the package installation step.

10.6 M1

For M1 clusters, upgrade will first update the streaming replicas one by one, then perform a *Switchover* from the primary to one of the replicas, update the primary, and switchover back to it again.

10.7 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.

If you did select specific versions, for example by using any of the `--xxx-package-version` options (e.g., postgres, bdr, pglogical) to

tpaexec configure, or by defining `xxx_package_version` variables in

`config.yml`, the upgrade will do nothing because the installed packages already satisfy the requested versions.

In this case, you must edit `config.yml`, remove the version settings, and re-run `tpaexec provision`. The update will then install the latest available packages. You can still update to a specific version by specifying versions on the command line as shown below:

```
$ tpaexec upgrade ~/clusters/speedy -vv \
-e postgres_package_version="2:11.6r2ndq1.6.13*" \
-e pglogical_package_version="2:3.6.11*" \
-e bdr_package_version="2:3.6.11*" \
```

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

tpaexec-configure.

It is your responsibility to ensure that the combination of Postgres, PGD, and pglogical package versions that you request are sensible. That is, they should work together, and there should be an upgrade path from what you have installed to the new versions.

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.

TPAEXEC SWITCHOVER

The `tpaexec switchover` command performs a controlled switchover between a primary and a replica in a [cluster that uses streaming replication](#). 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.)

11.1 Example

This command will make `replicaname` be the new primary in `~/clusters/speedy` :

```
$ tpaexec switchover ~/clusters/speedy replicaname
```

11.2 Architecture options

This command is applicable only to *M1 clusters*

that have a single writable primary instance and one or more read-only replicas.

For BDR-Always-ON clusters, use the [HAProxy server pool management commands](#) to perform maintenance on PGD instances.

BDR/HAPROXY SERVER POOL MANAGEMENT

The `tpaexec pool-disable-server` and `pool-enable-server` commands allow a PGD instance in a *BDR-Always-ON cluster* to be temporarily removed from the HAProxy active server pool for maintenance, and then added back afterwards.

These commands follow the same process as *rolling updates* 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.

12.1 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.

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 *streaming replication cluster*

, you can rehydrate the replicas first, then use *tpaexec switchover*

to convert the primary to a replica before rehydrating it. On

BDR-Always-ON clusters

, you can *remove each server from the haproxy server pool* 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 *tpaexec upgrade* command instead.

13.1 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.

13.2 Example

Let's assume you have an AWS cluster configuration in `~/clusters/night`.

13.2.1 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 change it using `awscli <#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.)

13.2.2 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.

13.3 Rehydrate in phases

In order to maintain cluster continuity, we recommend rehydrating the cluster in phases.

For example, in a [cluster that uses streaming replication](#) 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 [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.

13.4 Appendix

13.4.1 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.

- – title: Miscellaneous —

TPA AND ANSIBLE TOWER/ANSIBLE AUTOMATION CONTROLLER

TPA can be used with RedHat Ansible Automation Controller (formerly known as Ansible Tower) by running the configure and provision steps on a standalone machine, treating the cluster as a bare-metal one, and then importing the resulting inventory and other generated files into Tower.

A TPA installation on Tower instance, which includes its own virtual environment (tpa-venv), is then used for the deployment step in Tower.

This document describes the appropriate procedure for Ansible Tower version 3.8.

14.1 Preparing Tower for working with TPA

14.1.1 Setting up a TPA virtual environment

- Install `tpaexec` on your Tower server under `/opt/EDB/TPA`. This is the default location when TPA is installed from package.
- Run `tpaexec setup` which will create a virtual environment under `/opt/EDB/TPA/tpa-venv` and install the required packages.
- Add TPA directory (`/opt/EDB/TPA`) to the “CUSTOM VIRTUAL ENVIRONMENT PATHS” field in the System Settings page of Tower UI.

14.1.2 Creating the TPA_2Q_SUBSCRIPTION_TOKEN credential type

Create the custom credential type called “TPA_2Q_SUBSCRIPTION_TOKEN” as described below:

- Go to the Credentials Type page in Tower UI.
- Set the “NAME” field to “TPA_2Q_SUBSCRIPTION_TOKEN”.
- Paste this to “INPUT CONFIGURATION” field:

```
fields:
- id: tpa_2q_sub_token
  type: string
  label: TPA_2Q_SUBSCRIPTION_TOKEN
required:
- tpa_2q_subscription_token
```

- Paste this to “INJECTOR CONFIGURATION” field:

env:

```
TPA_2Q_SUBSCRIPTION_TOKEN: {{ tpa_2q_sub_token }}
```

- Click “SAVE” button.

14.2 Setting up a cluster

- Ensure the hosts you intend to use for your cluster are known to your Tower installation, that you have a Credential that can access them over SSH, and that they have `ansible_host`, `public_ip` and `private_ip` set.
- On a machine with `tpaexec` installed, prepare a file with a list of your hostnames, one per line:

```
mercury
venus
mars
jupiter
saturn
neptune
```

- Run `tpaexec configure`:

```
[tpa]$ tpaexec configure <clustername> \
    --platform bare \
    --use-ansible-tower https://aac.example.com/api \
    --tower-git-repository ssh://git@git.example.com/example \
    --hostnames-from <hostnamefile> \
    --architecture BDR-Always-ON \
    --layout bronze \
    --harp-consensus-protocol bdr \
    --postgresql 14
```

The API endpoint is currently accepted and added to `config.yml` but not used. The git repository will be used to import the cluster data into Tower; `tpaexec` will create its own branch in the repository for this cluster so it doesn’t matter what branches already exist in it. (This 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.

This will create a cluster directory named after your cluster.

- `config.yml` will now include the top-level dictionary `ansible_tower`, which will cause `tpaexec provision` to treat the cluster as a Tower cluster.
- Edit `config.yml` if you need to make any other changes.
- Run `tpaexec provision` to generate inventory and other related files.
- Add a Project in Tower with the git repository.
- Add the inventory from the project as an external source to your inventory.
- Create a Template in Tower specifying the TPA virtual environment, your inventory, and the newly created project. Also include any required credentials to reach your hosts, and a credential with a TPA subscription token (`TPA_2Q_SUBSCRIPTION_TOKEN`).
- Set one additional variable: `tpa_dir: /opt/EDB/TPA`
- Run a job based on the new Template to deploy your cluster.

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).

15.1 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-kfoodiiprztsyerriqbjqughhbemejgpc ; \  
    /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.

15.2 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.

15.3 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.

15.4 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.

15.5 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 `bpftrace/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.

15.6 Local privileges

The *installation instructions for TPA*

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.)

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 **tpaexecprovision**, 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.

16.1 Key conversion

Locateprivatekey

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**.

Savekeyas.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.

Keyconversion

Start **PuTTYgen** and under Parameters, select appropriate Type of key to generate:

For older versions of **PuTTYgen**, select **SSH-2RSA**; for recent versions select **RSA**

DonotselectSSH-1(RSA)

Now choose **Load** - in the box that says **PuTTYPrivateKeyFiles(*.ppk)** you will need to select **AllFiles(*.*)**

Select your **.pem** file and choose **Open**, then click **OK**.

Select **Saveprivatekey** 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**

16.2 ConfigurePuTTY

Start **PuTTY** and select **Session** from the **Category** window. In the **HostName** 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 ``tpaexecprovision`` utility.

In the Putty **Category** window, Select **Connection** , expand **SSH** and select **Auth**

For the panel marked **Privatekeyfileforauthentication** ,click** Browse **andselect-the.ppkfilethatwassavedabove,thenselect** Open**

In the Putty **Category** window, select **Session** again, enter a session name in **SavedSessions** , and **Save**

You should now be able to connect to the AWS host via PuTTY by selecting this saved session.

TROUBLESHOOTING

17.1 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):

```
[tpa]$ sudo rm -rf /opt/EDB/TPA/tpa-venv
[tpa]$ sudo /opt/EDB/TPA/bin/tpaexec setup
```

17.2 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.

```
[tpa]$ sudo ntpdate pool.ntp.org
```

17.3 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:

```
[tpa]$ tpaexec deploy <clustername> -v

- v      shows the results of modules
- vv     shows the files from which tasks come
- vvv    shows what commands are being executed on the target machines
- vvvv   enables connection debugging, what callbacks have been loaded
- vvvvv  shows some additional ssh configuration, filepath information
```

17.4 Cluster test

An easy way to smoketest an existing cluster is to run:

```
[tpa]$ 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:

```
[tpa]$ tpaexec test <clustername> --skip-tags pgbench
```

Tags in the test role are repmgr, postgres, barman, pgbench

Note that when specifying multiple tags, they should be comma delimited, with no spaces; for example:

```
[tpa]$ tpaexec test <clustername> --skip-tags barman,pgbench
```

17.5 TPA server test

To check the installation of the TPA server itself, run:

```
[tpa]$ tpaexec selftest
```

17.6 Skipping or including specific tags

When re-running a tpaexec provision or deploy after a failure, in the interests of time, it can sometimes be useful to miss out tasks by skipping specific tags. For example to miss out the repmgr tasks:

```
[tpa]$ tpaexec deploy <clustername> --skip-tags repmgr
```

To jump straight to re-run a particular task by specifying a tag—for example, to immediately run PGD tasks (the tag being bdr for legacy reasons):

```
[tpa]$ tpaexec deploy <bdr-clustername> --tags bdr
```

Notethatthisassumesthattheprevious tasksallcompletedsuccessfully.

To find all the tags for the relevant architecture that might be useful, run:

```
[tpa]$ grep -rs "tags:" /opt/EDB/TPA/architectures
```


REFERENCE

18.1 Running TPA in a Docker container

If you are using a system for which there are no [TPA packages](#) available, and it's difficult to run TPA after

Installing from source (for example, because it's not easy to obtain a working

Python 3.6+ 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 *deploy to Docker containers*. It's always preferable to run TPA directly if you can (even on MacOS X).

18.1.1 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/tpaexec.git
$ cd tpaexec
$ docker build -t tpa/tpaexec .
```

Double-check the created image:

```
$ docker image ls tpa/tpaexec
REPOSITORY    TAG       IMAGE ID       CREATED        SIZE
tpa/tpaexec    latest    e145cf8276fb   8 seconds ago  1.73GB
$ docker run --platform=linux/amd64 --rm tpa/tpaexec tpaexec info
# TPAexec v20.11-59-g85a62fe3 (branch: master)

tpaexec=/usr/local/bin/tpaexec
TPA_DIR=/opt/EDB/TPA
PYTHON=/opt/EDB/TPA/tpa-venv/bin/python3 (v3.7.3, venv)
TPA_VENV=/opt/EDB/TPA/tpa-venv
ANSIBLE=/opt/EDB/TPA/tpa-venv/bin/ansible (v2.8.15)
```

Create a TPA container and make your cluster configuration directories available inside the container:

```
$ docker run --platform=linux/amd64 --rm -v ~/clusters:/clusters \
-it tpa/tpaexec:latest
```

You can now run commands like `tpaexec provision /clusters/speedy` at the container prompt. (When you exit the shell, the container will be removed.)

If you want to provision Docker containers using TPA, you must also allow the container to access the Docker control socket on the host:

```
$ docker run --platform=linux/amd64 --rm -v ~/clusters:/clusters \
-v /var/run/docker.sock:/var/run/docker.sock \
-it tpa/tpaexec:latest
```

Run `docker ps` within the container to make sure that your connection to the host Docker daemon is working.

18.1.2 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 for Mac”](#) and launch Docker from the application menu.

18.2 Installing TPA from source

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

Note: EDB customers must [install TPA from packages](#) in order to receive EDB support for the software.

To run TPA from source, you must install all of the dependencies (e.g., Python 3.6+) that the packages would handle for you, or download the source and [run TPA in a Docker container](#). (Either way will work fine on Linux and macOS.)

18.2.1 Quickstart

First, you must install the various dependencies that would have been installed automatically along with the TPA packages. (You can use something other than `sudo` to run these commands as root, if you prefer.)

```
# Debian or Ubuntu

$ sudo apt-get install python3 python3-pip python3-venv \
git openvpn patch

# RedHat, Rocky or AlmaLinux (python3 for RHEL7, python36 for RHEL8)

$ sudo yum install python36 python3-pip \
epel-release git openvpn patch

# MacOS X

$ brew tap discoteq/discoteq
$ brew install python@3 openvpn flock coreutils gpatch git
```

Next, install TPA itself:

```
$ git clone ssh://git@github.com/EnterpriseDB/tpaexec.git
$ ./tpaexec/bin/tpaexec setup
$ ./tpaexec/bin/tpaexec selftest
```

18.2.2 Step-by-step

Install the various dependencies as described above.

If your system does not have Python 3.6+ packages, you can use `pyenv` to install a more recent Python in your home directory (see below), or you can *run TPA in a Docker container*.

Next, clone the TPA repository into, say, `~/tpaexec`. (It doesn't matter where you put it, but don't use `/opt/EDB/TPA` or `/opt/2ndQuadrant/TPA`, to avoid conflicts if you install the TPA packages in future.)

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

(If you're installing from source, please clone the repository instead of downloading an archive of the source.)

The remaining steps are the same as if you had installed the package.

```
# Add tpaexec to your PATH for convenience
# (Put this in your ~/.bashrc too)

$ export PATH=$PATH:$HOME/tpaexec/bin

$ tpaexec setup
$ tpaexec selftest
```

If the self-test completes without any errors, your TPA installation is ready for use.

18.2.3 Python 3.6+

TPA requires Python 3.6 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 3.9.0
Downloading Python-3.9.0.tar.xz...
- > https://www.python.org/ftp/python/3.9.0/Python-3.9.0.tar.xz
Installing Python-3.9.0...
Installed Python-3.9.0 to /home/ams/.pyenv/versions/3.9.0

$ pyenv local 3.9.0
$ pyenv version
3.9.0 (set by /home/ams/pyenv/.python-version)

$ pyenv which python3
/home/ams/.pyenv/versions/3.9.0/bin/python3
```

(continues on next page)

(continued from previous page)

```
$ python3 --version
3.9.0
```

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](#).

18.2.4 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"
```

18.3 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 *local-repo* and the required packages can be downloaded using the `tpaexec download-packages` command.

18.3.1 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 *using local-repo for distributing packages*

18.3.2 Downloading packages

On the internet connected machine, ensure that you have *docker installed* and run:

```
tpaexec download-packages cluster-dir --os <OS> --os-version <version>
```

See detailed description for the *package downloader*.

18.3.3 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.

18.3.4 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.

18.3.5 Updating in a disconnected environment

You can use the *upgrade* command to perform updates in an air-gapped environment. Prior to running this command you must run `download-packages` on the connected controller, copy the updated repository to the disconnected controller *and* then run `deploy` to copy these packages to the target instances.

18.4 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.

18.4.1 Debian x86

- Debian 11/bullseye is fully supported
- Debian 10/buster is fully supported
- Debian 9/stretch is a legacy distribution
- Debian 8/jessie is a legacy distribution

18.4.2 Ubuntu x86

- Ubuntu 22.04/jammy is fully supported
- Ubuntu 20.04/focal is fully supported
- Ubuntu 18.04/bionic is a legacy distribution
- Ubuntu 16.04/xenial is a legacy distribution

18.4.3 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 fully supported (python2 only)

18.4.4 SLES

- SLES 15.x is fully supported

18.4.5 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 [Docker containers](#).

18.5 TPA capabilities and supported software

- `:ref:Python requirements <Python environment>`
- `:ref:Supported distributions <Distribution support>`

18.5.1 Supported software

TPA can install and configure the following major components.

- Postgres 15, 14, 13, 12, 11
- EPAS (EDB Postgres Advanced Server) 15, 14, 13, 12
- 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)

18.6 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, and 5.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 [PGD documentation](#) for an authoritative description of the details.

18.6.1 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.

18.6.2 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). Set `bdr_package_version` to determine which exact package to install (e.g., '5.0*' to install the latest 5.0.x).

18.6.3 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.)

18.6.4 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.

18.6.5 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_node_group

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.

18.6.6 Miscellaneous notes

Hooks

TPA invokes the `bdr-node-pre-creation`, `bdr-post-group-creation`, and `bdr-pre-group-join` *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.

18.6.7 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.

18.7 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.

18.7.1 Barman configuration

On each barman server, a global configuration file will be created as `/etc/barman.conf` ; this contains default values for many barman configuration variables. For each postgres server being backed up, an additional Barman configuration file is created: to back up the server `one` , the file will be `/etc/barman.d/one.conf` , and the backups will be stored in `/var/lib/barman/one` . The file and directory names are taken from the backed-up instance's `backup_name` setting, defaulting to the instance's name.

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>	<code>/var/log/barman.log</code>
<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>	

18.7.2 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.

18.7.3 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.

18.8 Configuring EFM

TPA will install and configure EFM when `failover_manager` is set to `efm`.

Note that EFM is only available via EDB's package repositories and requires a valid subscription.

18.8.1 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.

See the [EFM documentation](#)

for more details on EFM configuration.

18.8.2 `efm_conf_settings`

You can use `efm_conf_settings` to set any parameters, whether recognised by TPA or not. Where needed, you need to quote the value exactly as it would appear in `efm.properties`:

```
cluster_vars:
  efm_conf_settings:
    standby.restart.delay: 1
    application.name: quarry
    reconfigure.num.sync: true
    reconfigure.num.sync.max: 1
    reconfigure.sync.primary: true
```

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`.

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 provide 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.

18.9 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 *Configuring 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
<code>haproxy_bind_address</code>	127.0.0.1	The address haproxy should bind to
<code>haproxy_port</code>	5432 (5444 for EPAS)	The TCP port haproxy should listen on
<code>haproxy_read_only_port</code>	5433 (5445 for EPAS)	TCP port for read-only load-balancer
<code>haproxy_backend_servers</code>	None	A list of Postgres instance names
<code>haproxy_maxconn</code>	<code>`max_connections`×0.9</code>	The maximum number of connections allowed per backend server; the default is derived from the backend's <code>max_connections</code> setting
<code>haproxy_peer_enabled</code>	True*	Add known haproxy hosts as <i>peer</i> list. *False` if <i>failover_manager</i> is harp or patroni.

18.9.1 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`.

18.9.2 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.

18.9.3 Example

```
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
```

18.10 Configuring HARP

TPA will install and configure HARP when `failover_manager` is set to `harp`, which is the default for BDR-Always-ON clusters.

18.10.1 Installing HARP

You must provide the `harp-manager` and `harp-proxy` packages. Please contact EDB to obtain access to these packages.

18.10.2 Configuring HARP

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.

continues on next page

Table 3 – continued from previous page

Variable	Default value	Description
<i>harp_lease_refresh_interval</i>	2000	Amount of time in milliseconds between re-freshes of the Lead Master lease.
<i>harp_dcs_reconnect_interval</i>	1000	The interval, measured in ms, between attempts that a disconnected node tries to reconnect to the DCS.
<i>harp_dcs_priority</i>	500	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>	false	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>	false	If HARP is unable to reach the DCS then fence the node.
<i>harp_maximum_lag</i>	1048576	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>	1048576	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>	lag_only	Whether CAMO queue state should be strictly enforced.
<i>harp_use_unix_sock</i>	false	Use unix domain socket for manager database access.
<i>harp_request_timeout</i>	250	Time in milliseconds to allow a query to the DCS to succeed.
<i>harp_watch_poll_interval</i>	500	Milliseconds to sleep between polling DCS. Only applies when <i>harp_consensus_protocol</i> is <i>bdr</i> .
<i>harp_proxy_timeout</i>	1	Builtin proxy connection timeout, in seconds, to Lead Master.
<i>harp_proxy_keepalive</i>	5	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>	75	Maximum number of client connections accepted by harp-proxy (<i>max_client_conn</i>)
<i>harp_ssl_password_command</i>	None	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>	10s	similar to dcs -> request_timeout, but for connection to the database itself.

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).

18.10.3 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.

18.10.4 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.

18.10.5 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.

18.10.6 Configuring the harp service

You can configure the following parameters for the harp service:

Variable	Default value	Description
<code>harp_manager_restart_on_failu</code>	<code>false</code>	If <code>true</code> , the <code>harp-manager</code> service is overridden so it's restarted on failure. The default is <code>false</code> to comply with the service installed by the <code>harp-manager</code> package.

18.10.7 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
<code>harp_http_options</code>	<code>enable: false`secure: false`host: <inventory_hostname>`port: 8080`probes: timeout: 10s`endpoint: "host=<proxy_name>port=<6432>dbname=<bdrdb>user=<username>"</code>	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.

18.11 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:

- `:ref: `Configure pg_hba.conf <pg_hba.conf>``

- `:ref: Configure postgresql.conf <postgresql.conf>`

Note that PEM is only available via EDB's package repositories and therefore requires a valid subscription.

18.11.1 Supported architectures

PEM is supported with M1 and BDR-Always-ON architectures via `--enable-pem` configuration command line option. You could 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.

Note that PEM server does not support pgextended for a backend yet.

18.11.2 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.

PEM server's web interface is configured to run on https and uses 443 port for the same. PEM's webserver configuration uses self-signed certificates.

Default login credentials for PEM server web interface use the postgres backend database user which is set to `postgres` for postgresql and `enterprisedb` for EPAS clusters by default. You could get the login password for the web interface by running `tpaexec show-password $clusterdir $user`.

18.11.3 Shared PEM server

Some deployments may want to use a single PEM sever for monitoring and managing multiple clusters in the organization. Shared pem server deployment within tpaexec is supported via `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
```

(continues on next page)

(continued from previous page)

```
private_ip: 10.33.15.102
vars:
  pem_shared: true
```

- 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
```

- 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.

- 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 pemserver'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.

18.11.4 Connecting to the PEM UI

PEM UI runs on https interface so you can connect with a running instance of PEM server via <https://\protect\T1\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.

18.12 Configuring pgbouncer

TPA will install and configure pgbouncer on instances whose `role` contains `pgbouncer` .

By default, pgbouncer listens for connections on port 6432 and forwards connections to `127.0.0.1:5432` (which may be either Postgres or

Configuring haproxy , depending on the architecture).

You can set the following variables on any pgbouncer instance.

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

18.12.1 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.

18.12.2 Authentication

PgBouncer will connect to Postgres as the `pgbouncer_auth_user` and execute the (already configured) `auth_query` to authenticate users.

18.12.3 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
```

18.13 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 architecture*

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`.)

18.13.1 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>	pgd-proxy 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. <code>enable_proxy_routing</code>

continues on next page

Table 8 – continued from previous page

Hash	Function	Description
<i>bdr_node_options</i>	<i>bdr.alter_node_option()</i>	routing configuration for individual PGD nodes

See the PGD documentation for more details.

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
      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

instances:
- Name: someproxy
  vars:
    pgd_proxy_options:
      listen_port: 9000
```

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>
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.

18.14 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
```

(continues on next page)

(continued from previous page)

```
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).

18.14.1 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 `postgres_users` and `postgres_databases` to create databases for replication, and use the

postgres-config-final

hook to populate the databases before pglogical is configured.

18.14.2 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.

18.14.3 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
```

(continues on next page)

(continued from previous page)

```

# 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`.)

18.14.4 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`.

18.14.5 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.

18.14.6 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.

18.15 Configuring 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`.

18.15.1 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.

18.15.2 repmgr on PGD instances

On PGD instances, `repmgr_failover` will be set to manual by default.

18.16 How TPA uses 2ndQuadrant and EDB repositories

This page explains the package sources from which TPA can download EDB (including 2ndQuadrant) software, how the source varies depending on the selected software, and how to configure access to each source.

Note that this page only describes the special configuration options and logic for EDB and 2ndQuadrant sources. Arbitrary *Configuring YUM repositories* or

Configuring APT repositories repositories can be added independently of the logic described here. Likewise, packages can be *downloaded in advance* and added to a *Creating and using a local repository* if preferred.

18.16.1 Package sources used by TPA

TPA downloads software from three package sources. Each of these sources provides multiple repositories. In some cases, the same software is available from more than one source.

- ``EDB Repos 2.0 <https://www.enterprisedb.com/repos/>`_`
- ``EDB Repos 1.0 <https://www.enterprisedb.com/repos/legacy>`_`
- ``2ndQuadrant Repos <https://techsupport.enterprisedb.com/customer_portal/sw/>`_`

By default, TPA will select sources and repositories automatically

based on the architecture and other options you have specified, so it is not generally necessary to change these. However, you will need to ensure that you have a valid subscription for all the sources used and that you have exported the token before running `tpaexec deploy` or the operation will fail.

Note: EDB is in the process of publishing all software through Repos 2.0, and will eventually remove the older repositories.

18.16.2 Authenticating with package sources

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](#).

To use [2ndQuadrant repositories](#), you must `export TPA_2Q_SUBSCRIPTION_TOKEN=xxx` before you run `tpaexec`. You can get your subscription token from the 2ndQuadrant Portal, under “Company info” in the left menu, then “Company”. Some repositories are available only by prior arrangement.

To use [EDB Repos 1.0](#) you must create a text file that contains your access credentials in the `username:password` format and run `export EDB_REPO_CREDENTIALS_FILE=/path/to/credentials/file` before you run `tpaexec`.

If you do not have an account for any of the sites listed, you can register for access at <https://www.enterprisedb.com/user/register?destination=/repository-access-request>

18.16.3 How sources are selected by default

If the PGD-Always-ON architecture is selected, repositories will be selected from EDB Repos 2.0 and all software will be sourced from these repositories.

If the M1 architecture is selected and no proprietary EDB software is selected, all packages will be sourced from PGDG. If any proprietary EDB software is selected, all packages will be sourced from EDB Repos 2.0.

For the BDR-Always-ON architecture, the default source is 2ndQuadrant, and the necessary repositories will be added from this source. In addition, the PGDG repositories will be used for community packages such as PostgreSQL and etcd as required. If EDB software not available in the 2ndQuadrant repos is required (e.g. EDB Advanced Server), repositories will be selected from EDB Repos 1.0.

18.16.4 Specifying EDB 2.0 repositories

To specify the complete list of repositories from EDB Repos 2.0 to install on each instance, set `edb_repositories` to a list of EDB repository names:

```
cluster_vars:
  edb_repositories:
    - enterprise
    - postgres_distributed
```

This example will configure the `enterprise` and `postgres_distributed` repositories, giving access to EPAS and PGD5 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.

If any EDB repositories are specified, any 2ndQuadrant repositories specified will be ignored and no EDB Repos 1.0 will be installed.

18.16.5 Specifying 2ndQuadrant repositories

To specify the complete list of 2ndQuadrant repositories to install on each instance in addition to the 2ndQuadrant public repository, set `tpa_2q_repositories` to a list of 2ndQuadrant repository names:

```
cluster_vars:
  tpa_2q_repositories:
    - products/pglogical3/release
    - products/bdr3/release
```

This example will install the `pglogical3` and `bdr3` release repositories. On Debian and Ubuntu systems, it will use the APT repository, and on RedHat systems, it will use the YUM repository.

The `dl/default/release` repository is always installed by default, unless you

- explicitly set `tpa_2q_repositories: []`, or
- have at least one entry in `edb_repositories`.

Either or the above will result in no 2ndQuadrant repositories being installed.

18.17 Configuring EDB Repos 2.0 repositories

This page explains how to configure EDB Repos 2.0 package repositories on any system.

For more details on the EDB and 2ndQuadrant package sources used by TPA see [this page](#).

To specify the complete list of repositories from EDB Repos 2.0 to install on each instance, set `edb_repositories` to a list of EDB repository names:

```
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 PGD5 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.

If any EDB repositories are specified, any 2ndQuadrant repositories specified will be ignored and no EDB Repos 1.0 will be installed.

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](#).

18.18 Configuring 2ndQuadrant repositories

This page explains how to configure 2ndQuadrant package repositories on any system.

For more details on the EDB and 2ndQuadrant package sources used by TPA see [this page](#).

To specify the complete list of 2ndQuadrant repositories to install on each instance in addition to the 2ndQuadrant public repository, set `tpa_2q_repositories` to a list of 2ndQuadrant repository names:

```
cluster_vars:
  tpa_2q_repositories:
    - products/pglogical3/release
    - products/bdr3/release
```

This example will install the pglogical3 and bdr3 release repositories. On Debian and Ubuntu systems, it will use the APT repository, and on RedHat systems, it will use the YUM repository. The 2ndQuadrant repositories are not available for SLES systems.

To use [2ndQuadrant repositories](#), you must export `TPA_2Q_SUBSCRIPTION_TOKEN=xxx` before you run `tpaexec`. You can get your subscription token from the 2ndQuadrant Portal, under “Company info” in the left menu, then “Company”. Some repositories are available only by prior arrangement.

The `dl/default/release` repository is always installed by default, unless you

- explicitly set `tpa_2q_repositories: []`, or
- have at least one entry in `edb_repositories`.

Either or the above will result in no 2ndQuadrant repositories being installed.

18.19 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: XXXXXXXX
      key_url: https://repo.example.com/path/to/XXXXXXXX.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.

18.20 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 *pre-deploy hook* 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`.

18.21 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 an air-gapped](air-gapped/) environment.

18.21.1 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 *details of the 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.

18.21.2 Creating a local repository manually

18.21.3 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/10`), 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.

18.21.4 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/10` , 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/buster  
# 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 .
```

18.21.5 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`.)

18.22 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
```

(continues on next page)

(continued from previous page)

```

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

SSH agent forwarding or an HTTPS username/password

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.

18.22.1 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 TPA packages*. For example

```

cluster_vars:
  packages:
    common:
      - golang-1.16

```

18.23 Git credentials

This page explains how to clone Git repositories that require authentication.

This may be required when you change `postgres_git_url` to

install Postgres from source or use ``install_from_source <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.

18.23.1 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.

18.23.2 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.

18.24 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 *extra_bashrc_lines* to set environment variables for the postgres user.

18.25 Python environment

TPA decides which Python interpreter to use based on the *distribution it detects* 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 10/buster	x	x (3.7)
Debian 9/stretch	x	x (3.5)
Debian 8/jessie	x	x (3.4)
Ubuntu 16.04/xenial	x	x (3.5)
Ubuntu 18.04/bionic	x	x (3.6)
Ubuntu 20.04/focal	x	x (3.8)
Ubuntu 22.04/jammy	x	x (3.10)
RHEL 7.x	x	x (3.6)
RHEL 8.x	x	x (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 Python 3 are not available, so TPA must use Python 2 instead. Debian 8 does not have the Python 3.5+ required to support Ansible.

You can decide for other distributions whether you prefer `python2` or `python3`, but the default for new clusters is `python3`.

18.25.1 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.

18.26 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
```

(continues on next page)

(continued from previous page)

```
...
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.

18.27 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
    - raid_device: /dev/md0
      device_name: /dev/xvdf
      volume_type: io2
      volume_size: 64
      raid_units: 2
      raid_level: 1
      iops: 5000
      vars:
```

(continues on next page)

(continued from previous page)

```

    volume_for: postgres_data
    encryption: luks
-   raid_device: /dev/md1
    device_name: /dev/xvdh
    ephemeral: ephemeral0
    raid_units: all
    vars:
        mountpoint: /mnt/scratch

```

In this example, the EC2 instance will end up with a 32GB EBS root volume, a 64GB RAID-1 volume comprising two provisioned-iops EBS volumes mounted as `/opt/postgres/data`, and a `/tmp/scratch` filesystem comprising all available instance-store (“ephemeral”) volumes, 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.

18.27.1 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
      volume_type: none

```

(continues on next page)

(continued from previous page)

```
- 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.)

18.27.2 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.

18.27.3 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.

18.27.4 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.

18.27.5 RAID arrays

On AWS EC2 instances, you can define RAID volumes:

```
instances:
- Name: one
  ...
  volumes:
  - raid_device: /dev/md0
    device_name: /dev/xvdf
    raid_units: 2
    raid_level: 1
    volume_type: gp2
```

(continues on next page)

(continued from previous page)

```

volume_size: 100
vars:
  volume_for: postgres_data

```

This example will attach 4×100GB EBS gp2 volumes (`/dev/xvd[f-i]`) and assemble them into a RAID-1 volume named `/dev/md0`. The handling of `volume_for` or `mountpoint` during deployment happens as with any other volume.

TPA does not currently support the creation and assembly of RAID arrays on other platforms, but you can use an existing array by adding an entry to `volumes` with `device_name: /dev/md0` or `/dev/mapper/xyz`. TPA will handle `mkfs` and `mount` as with any other block device.

18.27.6 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.

18.27.7 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: ...

```

(continues on next page)

(continued from previous page)

```

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.

18.28 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 *custom hook* .

18.29 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.)

18.30 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 instances* , 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 instances* . 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.

18.30.1 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.

18.31 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.

18.31.1 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 *SSH agent forwarding or an HTTPS username/password*

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.

18.31.2 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

install_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.

18.31.3 Package installation

There's a separate page about *installing Postgres and Postgres-related packages*

with `postgres_installation_method: pkg` (the default).

18.32 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:
```

(continues on next page)

(continued from previous page)

```
- 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 *extra_postgres_packages* instead.

18.32.1 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
```

18.32.2 Removing packages

You can specify a list of `unwanted_packages` that should be removed if they are installed.

```
unwanted_packages:
common:
- badpkg1
Ubuntu:
- badpkg2
```

18.33 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`.

18.33.1 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`.

18.34 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 *compiling and installing Postgres from source*.

18.35 SSL Certificates

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` will be installed via packages by default, but you can also install from a git branch or a local directory. See [configure-source.md](#) and [install_from_source.md](#) for more details.

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.

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

Testing.

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 it’s setup when both `--enable-pem` and `--enable-pg-backup-api` (or `config.yml` equivalent) are used.

18.36 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.

18.37 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:
        - 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 [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 [postgres_tablespace](#)).

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 [extra_postgres_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).

18.38 Creating Postgres tablespaces

To create Postgres tablespaces during deployment, define their names and locations in `postgres_tablespace` under `cluster_vars` or a particular instance's `vars` in `config.yml`.

If you [define volumes](#) with `volume_for: postgres_tablespace` set and a `tablespace_name` defined, they will be added as default entries to `postgres_tablespace` .

```
cluster_vars:
  postgres_tablespace:
    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_tablespace` 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 *postgres_users* will have been created by this time).

Streaming replicas must have the same `postgres_tablespace` volumes and `postgres_tablespace` setting as their upstream instance

You can set the default tablespace for a database in *postgres_databases*.

18.39 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`.

18.39.1 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.

18.39.2 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"`.

18.39.3 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"`.

18.39.4 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`.

18.39.5 Making changes by hand

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.

18.39.6 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`.

18.40 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.j2` in your cluster directory and set:

```
cluster_vars:
  postgres_hba_template: my_hba.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.

18.41 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 %}
```

18.42 Configuring .pgpass

TPA will create `~postgres/.pgpass` by default with the passwords for `postgres` and `repmgr` in it, for use between cluster instances. You can set `pgpass_users` to create entries for a different list of users.

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
```

18.43 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 *Postgres users in the database*.

18.43.1 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` in 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.

18.43.2 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 .

18.43.3 TLS certificates

By default, TPA will generate a private key and a self-signed TLS certificate for use within the cluster. This is sufficient to ensure that traffic between clients and server is encrypted in transit. Should you wish to use your own certificate signing infrastructure you may replace these after deployment is complete, or replace them during deployment using a *Hooks* .

18.43.4 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.

18.44 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.

18.44.1 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
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
)>tkc}}kly4&epaJ?;NJ:luT{C7D*<p
$ tpaexec store-password ~/clusters/speedy example
Password:
$ tpaexec show-password ~/clusters/speedy example
terrible insecure password
$
```

If you don't want the user to have a password at all, you can set `generate_password: false`.

18.45 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.

18.45.1 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

Provisioning the cluster. You will get more logs if you also *Deployment* the cluster.

18.45.2 Quickstart

```
[tpa]$ tpaexec archive-logs ~/clusters/speedy

PLAY [Prepare local host archive] *****
  ↳ *****

TASK [Collect facts] *****
  ↳ *****
ok: [localhost]

TASK [Set time stamp] *****
```

(continues on next page)

(continued from previous page)

```

↪ ****  ****  ***
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.

18.45.3 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
|   |-- journal.log
|   |-- ks-script-ipdkisn0.log
|   |-- ks-script-jr03uzns.log

```

(continues on next page)

(continued from previous page)

```
| |-- 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
```

18.46 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.

18.46.1 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 [configuring a cluster](#)

.

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`.

18.46.2 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 [recommendations for installing to an air-gapped environment](#) . A *local-repo* will be detected and used automatically by TPA.

18.46.3 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
```

18.47 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.

18.47.1 Quickstart

- Create `commands/mycmd.yml` within your cluster directory
- Run `tpaexec mycmd /path/to/cluster`

18.47.2 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
```

18.48 TPA custom tests

You can easily define in-depth tests specific to your environment and application to augment TPA's *builtin tests*.

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.

18.48.1 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.)

18.48.2 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.

18.48.3 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.)

18.49 Shared instances

18.50 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.

18.51 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
```

If deploying to RedHat you must also add the PGDG repository to your yum repository list in config.yml:

```
cluster_vars:
  yum_repository_list:
    - PGDG
```

TPA configure will add 3 etcd nodes and 2 haproxy nodes. Etcd is used for the Distributed Configuration Store (DCS). Patroni supports other DCS backends, but they are not currently supported by EDB or TPA.

TPA uses Patroni's feature of converting an existing PostgreSQL standalone server. This allows for TPA to initialise and manage configuration. Once a single PostgreSQL server and database has been created, Patroni will create replicas and configure replication. 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.

These configuration variables can be used to control certain behaviours in the deployment of Patroni in TPA.

Variable	Default value	Description
<i>patroni_super_user</i>	postgres	User to create in postgres for superuser role.
<i>patroni_replication_user</i>	replicator	Username to create in postgres for replication role.
<i>patroni_restapi_user</i>	patroni	Username to configure for the patroni REST API.
<i>patroni_rewind_user</i>	rewind	Username to create in postgres for pg_rewind function.
<i>patroni_installation_method</i>	pkg	Install patroni from packages or source (e.g. git repo or local source directory if docker).
<i>patroni_ssl_enabled</i>	no	Whether to enable SSL for REST API and ctl connection. Will use the cluster SSL cert and CA if available.
<i>patroni_rewind_enabled</i>	yes	Whether to enable postgres rewind, creates a user defined by patroni_rewind_user 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_listen_port</i>	8008	REST API TCP port number

continues on next page

Table 10 – continued from previous page

Variable	Default value	Description
<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 tpaexec-configure for possible values for <i>/replica?lag</i>

18.51.1 Patroni configuration file settings

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.md](#) . 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
```

TPA provides these minimal set of tools for managing Patroni clusters.

18.51.2 Status

To see the current status of the TPA cluster according to Patroni run

```
tpaexec status cluster_name
```

18.51.3 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.

18.52 tpaexec info

You can use the `info` command to output information about the TPA installation. Providing this information is valuable for troubleshooting.

18.52.1 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

18.52.2 Example Output

The `tpaexec info` command outputs the following:

```
# TPAexec 23.18.18

tpaexec=./tpaexec
TPA_DIR=/opt/EDB/TPA
PYTHON=/opt/EDB/TPA/tpa-venv/bin/python3 (v3.7.3, venv)
TPA_VENV=/opt/EDB/TPA/tpa-venv
ANSIBLE=/opt/EDB/TPA/tpa-venv/bin/ansible (v2.9.27)
Validation: e05e5302cd357b8ddbb042b7591bf66dfa283213ccbe5073b2cff3c783be1310 [OK]
```

18.53 tpaexec reconfigure

The `tpaexec reconfigure` command reads `config.yml` and generates a revised version of it that changes the cluster from one architecture to another. *tpaexec upgrade* may then be invoked to make the required changes on the instances that make up the cluster.

18.53.1 Arguments

As with other `tpaexec` commands, the cluster directory must always be given.

The following arguments control the contents of the new `config.yml`:

- `--architecture <architecture>` (required) The new architecture for the cluster. At present the only supported architecture is `PGD-Always-ON`
- `--pgd-proxy-routing <global|local>` (required) How PGD-Proxy is to route connections. See [the PGD-Always-ON documentation](#) 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.

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`.

18.53.2 Sample invocation

```
$ tpaexec reconfigure ~/clusters/speedy\
--architecture PGD-Always-ON\
--pgd-proxy-routing local
```