

Introduction

Contents

Crunchy Data Container Suite	1
Overview	2
User Guide	8
Running the Examples	9
pgAdmin4 with TLS	26
pgBackRest Examples	27
pgBaseBackup Examples	38
PostgreSQL Container Example	54
Requirements	57
Project Environment	57
Installation	57
Configuring Namespace and Permissions	60
Storage Configuration	61

Crunchy Data Container Suite

Crunchy Container Suite is a collection of container images for PostgreSQL databases and various tools to manage them. The container images provided aim to be highly configurable, easy to deploy and reliable.

The suite is cloud agnostic and certified for production deployment with leading container platform technologies. Crunchy Container Suite supports all major public clouds and enables hybrid cloud deployments.

Crunchy Container Suite includes:

- Compatibility with Docker, Kubernetes and OpenShift
- Cloud-agnostic: build your own database-as-a-service in any public, private, or hybrid cloud
- No forks: 100% open source, native PostgreSQL
- Backup, restore, & disaster recovery for terabytes of data
- Graphical database administration tools for simple point-and-click database management
- Open source database monitoring and graphing tools to help analyze and create administration strategies
- PostGIS, robust open source GIS functionality, included by default
- Access to certified builds for Crunchy Certified PostgreSQL
- No proprietary software licensing fees

Why Crunchy Containers?

Enterprise production PostgreSQL deployments require more than just a database. Advanced capabilities like high availability, high-performant disaster recovery for terabytes of data, and monitoring solutions are requirements for your enterprise database cluster.

By providing the necessary microservices, including containers for scaling, high-availability, disaster recovery, monitoring, and more, Crunchy Container Suite will meet your production compliance, security, and performance requirements and give you a trusted open source database experience.

Elastic PostgreSQL

From creating uniformly managed, cloud-native production deployments to allowing your engineering team to provision databases that meet your compliance requirements, Crunchy Container Suite gives your organization the flexibility to deploy your own personalized database-as-a-service tailored to your needs.

Open Source For Enterprise

Crunchy Container Suite comes with essential open source tools for PostgreSQL management at scale, and lets you use powerful extensions like geospatial management with PostGIS.

Compliance At Scale

Deploy Crunchy Certified PostgreSQL with Crunchy Container Suite to harness the security of a Common Criteria EAL 2+ certified database on trusted platforms such as Red Hat OpenShift or Pivotal Container Service.

Overview

The following provides a high level overview of each of the container images.

CentOS vs RHEL Images

The Crunchy Container suite provides two different OS images: `centos7` and `rhe17`. These images are identical except for the packages used by `yum` to install the software.

The `centos7` images, `yum` is configured to use PostgreSQL RPM Building Project.

The `rhe17` images use Crunchy Certified RPMs and are only available to active Crunchy Data customers.

Database Images

Crunchy Container Suite provides two types of PostgreSQL database images:

- Crunchy PostgreSQL
- Crunchy PostGIS

Supported major versions of these images are:

- 9.5
- 9.6
- 10
- 11

Crunchy PostgreSQL

Crunchy PostgreSQL is an unmodified deployment of the PostgreSQL relational database. It supports the following features:

- Asynchronous and synchronous replication
- Mounting custom configuration files such as `pg_hba.conf`, `postgresql.conf` and `setup.sql`
- Can be configured to use SSL authentication
- Logging to container logs
- Dedicated users for: administration, monitoring, connection pooler authentication, replication and user applications.
- pgBackRest backups built into the container
- Archiving WAL to dedicated volume mounts
- [Extensions available](#) in the PostgreSQL contrib module.
- Enhanced audit logging from the pgAudit extension
- Enhanced database statistics from the `pg_stat_statements` extensions

Crunchy PostgreSQL PostGIS

The Crunchy PostgreSQL PostGIS mirrors all the features of the Crunchy PostgreSQL image but additionally provides the following geospatial extensions:

- PostGIS
- PostGIS Topology
- PostGIS Tiger Geocoder
- FuzzyStrMatch
- PLR

Backup and Restoration Images

Crunchy Container Suite provides two types of backup images:

- Physical - backups of the files that comprise the database
- Logical - an export of the SQL that recreates the database

Physical backup and restoration tools included in the Crunchy Container suite are:

- [pgBackRest](#) - included and configurable in the Crunchy PostgreSQL images
- [pg_basebackup](#) - provided by the Crunchy Backup image

Logical backup and restoration tools are:

- [pg_dump](#) - provided by the Crunchy pgDump image
- [pg_restore](#) - provided by the Crunchy pgRestore image

Crunchy Backup

The Crunchy Backup image allows users to create [pg_basebackup](#) physical backups. The backups created by Crunchy Backup can be mounted to the Crunchy PostgreSQL container to restore databases.

Crunchy BackRest Restore

The Crunchy BackRest Restore image restores a PostgreSQL database from pgBackRest physical backups. This image supports the following types of restores:

- Full - all database cluster files are restored and PostgreSQL replays Write Ahead Logs (WAL) to the latest point in time. Requires an empty data directory.
- Delta - missing files for the database cluster are restored and PostgreSQL replays Write Ahead Logs (WAL) to the latest point in time.
- PITR - missing files for the database cluster are restored and PostgreSQL replays Write Ahead Logs (WAL) to a specific point in time.

Visit the official pgBackRest website for more information: <https://pgbackrest.org/>

Crunchy pgDump

The Crunchy pgDump image creates a logical backup of the database using the `pg_dump` tool. It supports the following features:

- `pg_dump` individual databases
- `pg_dump` all databases
- various formats of backups: plain (SQL), custom (compressed archive), directory (directory with one file for each table and blob being dumped with a table of contents) and tar (uncompressed tar archive)
- Logical backups of database sections such as: DDL, data only, indexes, schema

Crunchy pgRestore

The Crunchy pgRestore image allows users to restore a PostgreSQL database from `pg_dump` logical backups using the `pg_restore` tool.

Administration

The following images can be used to administer and maintain Crunchy PostgreSQL database containers.

Crunchy pgAdmin4

The Crunchy pgAdmin4 images allows users to administer their Crunchy PostgreSQL containers via a graphical user interface web application.




Figure 1: pgadmin4

Visit the official pgAdmin4 website for more information: <https://www.pgadmin.org/>

Crunchy Scheduler

The Crunchy Scheduler image provides a cronlike microservice for automating `pg_basebackup` and `pgBackRest` backups within a single Kubernetes namespace.

The scheduler watches Kubernetes for config maps with the label `crunchy-scheduler=true`. If found the scheduler parses a JSON object contained in the config map and converts it into an scheduled task.

Crunchy Upgrade

The Crunchy Upgrade image allows users to perform major upgrades of their Crunchy PostgreSQL containers. The following upgrade versions of PostgreSQL are available:

- 9.5
- 9.6
- 10
- 11

Performance and Monitoring

The following images can be used to understand how Crunchy PostgreSQL containers are performing over time using tools such as Grafana, Prometheus and pgBadger.

Crunchy Collect

The Crunchy Collect image exports metric data of Crunchy PostgreSQL containers which can is scraped and stored by Crunchy Prometheus timeseries database via a web API.

Crunchy Collect contains the following exporters:

- [Node Exporter](#) - hardware and OS metrics
- [PostgreSQL Exporter](#) - postgres specific metrics

This image also contains custom PostgreSQL queries for additional metrics provided by [Crunchy pgMonitor](#).

Crunchy Grafana

The Crunchy Grafana image provides a web interface for users to explore metric data gathered and stored by Prometheus. Crunchy Grafana comes with the following features:

- Premade dashboards tuned for PostgreSQL metrics
- Automatic datasource registration
- Automatic administrator user setup

Visit the official Grafana website for more information: <https://grafana.com>

Crunchy Prometheus

The Crunchy Prometheus image provides a time series databases for storing metric data gathered from Crunchy PostgreSQL containers. Metrics can be explored via queries in the Prometheus graphical user interface and visualized using Crunchy Grafana. Crunchy Prometheus supports the following features:

- Auto discovers metric exporters in Kubernetes by searching for pods with the label `crunchy-collect=true`
- Relabels metrics metadata for easier Crunchy Grafana integration

Visit the official Prometheus website for more information: <https://prometheus.io>




Figure 2: grafana

Crunchy pgBadger

The Crunchy pgBadger image provides a tool that parses PostgreSQL logs and generates an in-depth statistical report. Crunchy pgBadger reports include:

- Connections
- Sessions
- Checkpoints
- Vacuum
- Locks
- Queries

Additionally Crunchy pgBadger can be configured to store reports for analysis over time.




Figure 3: pgbadger

Visit the official pgBadger website for more information: <https://pgbadger.darold.net/>

Connection Pooling and Logical Routers

Crunchy pgBouncer

The Crunchy pgBouncer image provides a lightweight PostgreSQL connection pooler. Using pgBouncer, users can lower overhead of opening new connections and control traffic to their PostgreSQL databases. Crunchy pgBouncer supports the following features:

- Connection pooling
- Drain, Pause, Stop connections to Crunchy PostgreSQL containers
- Dedicated pgBouncer user for authentication queries
- Dynamic user authentication

Visit the official pgBouncer website for more information: <https://pgbouncer.github.io>

Crunchy pgPool II

The Crunchy pgPool image provides a logical router and connection pooler for Crunchy PostgreSQL containers. pgPool examines SQL queries and redirects write queries to the primary and read queries to replicas. This allows users to setup a single entrypoint for their applications without requiring knowledge of read replicas. Additionally pgPool provides connection pooling to lower overhead of opening new connections to Crunchy PostgreSQL containers.

Visit the official pgPool II website for more information: <http://www.pgpool.net>

Supported Platforms

Crunchy Container Suite supports the following platforms:

- *Docker 1.13+*
- *Kubernetes 1.8+*
- *OpenShift Container Platform 3.11*

User Guide

Overview

This guide is intended to get you up and running with the Crunchy Container Suite, and therefore provides guidance for deploying the Crunchy Container Suite within your own environment. This includes guidance for standing-up and configuring your environment in order to run Crunchy Containers examples that can be found in the next section.

Please see the following sections in order to properly setup and configure your environment for the Crunchy Container Suite (*please feel free to skip any sections that have already been completed within your environment*):

1. [Platform Installation](#)
2. [Crunchy Container Suite Installation](#)
3. [Storage Configuration](#)
4. [Example Guidance](#)
5. [Crunchy Container Suite Examples](#)

Once your environment has been configured according to instructions provided above, you will be able to run the Crunchy Container Suite examples. These examples will demonstrate the various capabilities provided by the Crunchy Container Suite, including how to properly configure and deploy the various containers within the suite, and then utilize the features and services provided by those containers. The examples therefore demonstrate how the Crunchy Container Suite can be utilized to effectively deploy a PostgreSQL database cluster within your own environment, that meets your specific needs and contains the PostgreSQL features and services that you require.

Platform Installation

In order to run the examples and deploy various containers within the Crunchy Container Suite, you will first need access to an environment containing one of the following supported platforms:

- Docker 1.13+ (<https://www.docker.com/>)
- Kubernetes 1.8+ (<https://kubernetes.io/>)
- OpenShift Container Platform 3.11 (<https://www.openshift.com/products/container-platform/>)

Links to the official website for each of these platform are provided above. Please consult the official documentation for instructions on how to install and configure these platforms in your environment.

Crunchy Container Suite Installation

Once you have access to an environment containing one of the supported platforms, it is then necessary to properly configure that environment in order to run the examples, and therefore deploy the various containers included in the Crunchy Container Suite. This can be done by following the Crunchy Container Suite [Installation Guide](#).

Please note that as indicated within the [Installation Guide](#), certain steps may require administrative access and/or privileges. Therefore, please work with your local System Administrator(s) as needed to setup and configure your environment according to the steps defined within this guide. Additionally, certain steps are only applicable to certain platforms and/or environments, so please be sure to follow all instructions that are applicable to your target environment.

Storage Configuration

Once you have completed all applicable steps in the [Installation Guide](#), you can then proceed with configuring storage in your environment. The specific forms of storage supported by the Crunchy Containers Suite, as well as instructions for configuring and enabling those forms of storage, can be found in the [Storage Configuration](#) guide. Therefore, please review and follow steps in the [Storage Configuration](#) guide in order to properly configure storage in your environment according to your specific storage needs.

Example Guidance

With the [Installation Guide](#) and [Storage Configuration](#) complete, you are almost ready to run the examples. However, prior to doing so it is recommended that you first review the documentation for [Running the Examples](#), which describes various conventions utilized in the examples, while also providing any other information, resources and guidance relevant to successfully running the Crunchy Container Suite examples in your environment. The documentation for running the examples can be found [here](#).

Running the Examples

The Kubernetes and OpenShift examples in this guide have been designed using single-node Kubernetes/OCP clusters whose host machines provide any required supporting infrastructure or services (e.g. local HostPath storage or access to an NFS share). Therefore, for the best results when running these examples, it is recommended that you utilize a single-node architecture as well.

Additionally, the examples located in the **kube** directory work on both Kubernetes and OpenShift. Please ensure the `CCP_CLI` environment variable is set to the correct binary for your environment, as shown below:

```
# Kubernetes
export CCP_CLI=kubectl

# OpenShift
export CCP_CLI=oc
```

***NOTE:** Set the `CCP_CLI` environment variable in `.bashrc` to ensure the examples will work properly in your environment*

Example Conventions

The examples provided in Crunchy Container Suite are simple examples that are meant to demonstrate key Crunchy Container Suite features. These examples can be used to build more production level deployments as dictated by user requirements specific to their operating environments.

The examples generally follow these conventions: - There is a **run.sh** script that you will execute to start the example - There is a **cleanup.sh** script that you will execute to shutdown and cleanup the example - Each example will create resources such as Secrets, ConfigMaps, Services, and PersistentVolumeClaims, all which follow a naming convention of `<example name>-<optional description suffix>`. For example, an example called **primary** might have a PersistentVolumeClaim called **primary-pgconf** to describe the purpose of that particular PVC. - The folder names for each example give a clue as to which Container Suite feature it demonstrates. For instance, the `examples/kube/pgaudit` example demonstrates how to enable the **pg_audit** capability in the **crunchy-postgres** container.

Helpful Resources

Here are some useful resources for finding the right commands to troubleshoot and modify containers in the various environments shown in this guide:

- [Docker Cheat Sheet](#)
- [Kubectl Cheat Sheet](#)
- [OpenShift Cheat Sheet](#)
- [Helm Cheat Sheet](#)

Crunchy Container Suite Examples

Now that your environment has been properly configured for the Crunchy Container Suite and you have reviewed the guidance for running the examples, you are ready to run the Crunchy Container Suite examples. Therefore, please proceed to the next section in order to find the examples that can now be run in your environment. — title: “Troubleshooting” date: draft: false weight: 8 —

Kubernetes

[Troubleshooting kubeadm](#)

509 Certificate Errors

If you see `Unable to connect to the server: x509: certificate has expired or is not yet valid`, try resetting ntp. This generally indicates that the date/time is not set on local system correctly.

If you see `Unable to connect to the server: x509: certificate signed by unknown authority` (possibly because of "crypto/rsa: verification error" while trying to verify candidate authority certificate "kubernetes"), try running these commands as a regular user:

```
mv $HOME/.kube $HOME/.kube.bak
mkdir -p $HOME/.kube
sudo cp -i /etc/kubernetes/admin.conf $HOME/.kube/config
sudo chown $(id -u):$(id -g) $HOME/.kube/config
```

gcloud Errors

If you see the error `ERROR: (gcloud.container.clusters.get-credentials)Unable to create private file [/etc/kubernetes/adm [Errno 1] Operation not permitted: '/etc/kubernetes/admin.conf'`, create a backup of admin.conf and delete the admin.conf before attempting to reconnect to the cluster.

gcloud Authentication Example

The commands used to authenticate to gcloud are the following:

```
gcloud auth login
gcloud config set project <your gcloud project>
gcloud auth configure-docker
```

If you see gcloud authentication errors, execute `gcloud config list` then re-authenticate using the commands from above. Finally, rerun `gcloud config list` - the results should show different values if authentication was indeed the issue.

OpenShift Container Platform

[Troubleshooting OpenShift Container Platform: Basics](#)

The `crunchy-backrest-restore` container executes the `pgBackRest` utility, allowing FULL and DELTA restore capability. See the [pgBackRest](#) guide for more details.

Features

The following features are supported and required by the `crunchy-backrest-restore` container:

- Mounted `pgbackrest.conf` configuration file via the `/pgconf` volume
- Mounted `/backrestrepo` for access to `pgBackRest` archives

Packages

The `crunchy-backrest-restore` Docker image contains the following packages (versions vary depending on PostgreSQL version):

- PostgreSQL (11.1, 10.6, 9.6.11 and 9.5.15)
- [pgBackRest](#) (2.x)
- CentOS7 - publicly available
- RHEL7 - customers only

Environment Variables

Required

Name	Default	Description
PGBACKREST_STANZA	None	Must be set to the desired stanza for restore.

Optional

Name	Default	Description
PGBACKREST_DELTA	None	Enables pgBackRest delta restore mode. Used when a user needs to restore to a volume.
PGBACKREST_TARGET	None	PostgreSQL timestamp used when restoring up to a point in time. Required for Point In Time.
PGBACKREST_PG1_PATH	None	Path where PostgreSQL data directory can be found. This variable can also be used to specify the path to the backup.
BACKREST_CUSTOM_OPTS	None	Custom pgBackRest options can be added here to customize pgBackRest restores.
CRUNCHY_DEBUG	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal secrets in logs.

The crunchy-backup container executes a full backup against another database container using the standard `pg_basebackup` utility that is included with PostgreSQL.

Features

The following features are supported by the `crunchy-backup` container:

- Backup and restoration from: `pg_basebackup`

Packages

The crunchy-backup Docker image contains the following packages (versions vary depending on PostgreSQL version):

- PostgreSQL (11.1, 10.6, 9.6.11 and 9.5.15)
- CentOS7 - publicly available
- RHEL7 - customers only

Environment Variables

Required

Name	Default	Description
BACKUP_LABEL	crunchy-backup	The label for the backup.
BACKUP_HOST	None	Name of the database the backup is being performed on.
BACKUP_USER	None	Username for the PostgreSQL role being used.
BACKUP_PASS	None	Password for the PostgreSQL role being used.
BACKUP_PORT	5432	Database port used to do the backup.

Optional

Name	Default	Description
CRUNCHY_DEBUG	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal secrets in logs.
BACKUP_OPTS	None	Optional parameters to pass to <code>pg_basebackup</code> .

Volumes

Name	Description
<code>/backup</code>	Volume used by the <code>pg_basebackup</code> backup tool to store physical backups.
<code>/pgdata</code>	Volume used to store the data directory contents for the PostgreSQL database.

Backup Location

Backups are stored in a mounted backup volume location, using the database host name plus `-backups` as a sub-directory, then followed by a unique backup directory based upon a date/timestamp. It is left to the user to perform database backup archives in this current version of the container. This backup location is referenced when performing a database restore.

The `crunchy-collect` container provides real time metrics about the PostgreSQL database via an API. These metrics are scraped and stored by a [Prometheus](#) time-series database and are then graphed and visualized through the open source data visualizer [Grafana](#).

The `crunchy-collect` container uses [pgMonitor](#) for advanced metric collection. It is required that the `crunchy-postgres` container has the `PGMONITOR_PASSWORD` environment variable to create the appropriate user (`ccp_monitoring`) to collect metrics.

Custom queries to collect metrics can be specified by the user. By mounting a `queries.yml` file to `/conf` on the container, additional metrics can be specified for the API to collect. For an example of a `queries.yml` file, see [here](#)

Packages

The `crunchy-collect` Docker image contains the following packages (versions vary depending on PostgreSQL version):

- PostgreSQL (11.1, 10.6, 9.6.11 and 9.5.15)
- CentOS7 - publicly available
- RHEL7 - customers only
- [PostgreSQL Exporter](#)
- [Node Exporter](#)

Environment Variables

Required

Name	Default	Description
<code>DATA_SOURCE_NAME</code>	None	The URL for the PostgreSQL server's data source name. This is <i>required</i> to be in the form of

Optional

Name	Default	Description
<code>DISABLE_NODE_EXPORTER</code>	false	Set this to true to disable node exporter in the collect container.
<code>CRUNCHY_DEBUG</code>	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal secrets in logs

Visual dashboards are created from the collected and stored data that `crunchy-collect` and `crunchy-prometheus` provide for the `crunchy-grafana` container, which hosts an open source web-based graphing dashboard called [Grafana](#).

Grafana is a platform which can then apply the defined metrics and visualize information through various tools. It is extremely flexible with a powerful query and transformation language, producing beautiful and easily understandable graphics to analyze and monitor your data.

By default, `crunchy-grafana` will register the [Prometheus](#) datasource within Grafana and import a pre-made dashboard for PostgreSQL monitoring.

The `crunchy-grafana` container must be able to reach the `crunchy-prometheus` container.

Users must specify an administrator username and password to provide basic authentication for the web frontend. Additionally, the Prometheus hostname and port number are required. If Prometheus uses basic authentication, users must specify the username and password to access Prometheus via environment variables.

A user may define a custom `defaults.ini` file and mount to `/conf` for custom configuration. For configuration examples, see [here](#).

The following port is exposed by the `crunchy-grafana` container:

- `crunchy-grafana:3000` - the Grafana web user interface

Packages

The `crunchy-grafana` Docker image contains the following packages:

- [Grafana](#)
- CentOS7 - publicly available
- RHEL7 - customers only

Environment Variables

Required

Name	Default	Description
<code>ADMIN_USER</code>	None	Specifies the administrator user to be used when logging into the web frontend.
<code>ADMIN_PASS</code>	None	Specifies the administrator password to be used when logging into the web frontend.
<code>PROM_HOST</code>	None	Specifies the Prometheus container hostname for auto registering the Prometheus datasource.
<code>PROM_PORT</code>	None	Specifies the Prometheus container port for auto registering the Prometheus datasource.

Optional

Name	Default	Description
<code>PROM_USER</code>	5s	Specifies the Prometheus username, if one is required.
<code>PROM_PASS</code>	5s	Specifies the Prometheus password, if one is required.
<code>CRUNCHY_DEBUG</code>	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal secrets in logs.

The `crunchy-pgadmin4` container executes the [pgAdmin4](#) web application.

`pgAdmin4` provides a web user interface to PostgreSQL databases. A sample screenshot is below:

Features

The following features are supported by the `crunchy-pgadmin4` container:

- Expose port (5050 by default) which is the web server port.
- Mount a certificate and key to the `/certs` directory and set `ENABLE_TLS` to true to activate HTTPS mode.
- Set username and password for login via environment variables.

Restrictions

- An emptyDir, with write access, must be mounted to the `/run/httpd` directory in OpenShift.

Packages

The `crunchy-pgadmin4` Docker image contains the following packages (versions vary depending on PostgreSQL version):




Figure 4: pgAdmin4

- PostgreSQL (11.1, 10.6, 9.6.11 and 9.5.15)
- [pgAdmin4](#)
- CentOS7 - publicly available
- RHEL7 - customers only

Environment Variables

Required

Name	Default	Description
PGADMIN_SETUP_EMAIL	None	Set this value to the email address used for pgAdmin4 login.
PGADMIN_SETUP_PASSWORD	None	Set this value to a password used for pgAdmin4 login. This should be a strong password.
SERVER_PORT	5050	Set this value to a change the port pgAdmin4 listens on.
ENABLE_TLS	FALSE	Set this value to true to enable HTTPS on the pgAdmin4 container. This requires a certificate.

Optional

Name	Default	Description
CRUNCHY_DEBUG	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal secrets in logs.

The crunchy-pgbadger container executes the [pgBadger](#) utility, which generates a PostgreSQL log analysis report using a small HTTP server running on the container. This log report can be accessed through the URL `http://<>:10000/api/badgergenerate`.

Features

The following features are supported by the crunchy-pgbouncer container:

- crunchy-pgbouncer uses `auth_query` to authenticate users. This requires only the `pgbouncer` username and password in `users.txt`. Automatically generated from environment variables.
- Mount a custom `users.txt` and `pgbouncer.ini` configurations for advanced usage.
- Tune pooling parameters via environment variables.
- Connect to the administration database in pgBouncer to view statistics of the target databases.

Packages

The crunchy-pgbouncer Docker image contains the following packages (versions vary depending on PostgreSQL version):

- PostgreSQL (11.1, 10.6, 9.6.11 and 9.5.15)
- [pgBouncer](#)
- CentOS7 - publicly available
- RHEL7 - customers only

Restrictions

- OpenShift: If custom configurations aren't being mounted, an `emptydir` volume is required to be mounted at `/pgconf`.
- Superusers cannot connect through the connection pooler.

Environment Variables

Required

Name	Default	Description
<code>PGBOUNCER_PASSWORD</code>	None	The password of the pgBouncer role in PostgreSQL. Must be also set on the primary database.
<code>PG_SERVICE</code>	None	The hostname of the database service.

Optional

Name	Default	Description
<code>DEFAULT_POOL_SIZE</code>	20	How many server connections to allow per user/database pair.
<code>MAX_CLIENT_CONN</code>	100	Maximum number of client connections allowed.
<code>MAX_DB_CONNECTIONS</code>	Unlimited	Do not allow more than this many connections per-database.
<code>MIN_POOL_SIZE</code>	0	Adds more server connections to pool if below this number.
<code>POOL_MODE</code>	Session	When a server connection can be reused by other clients. Possible values are <code>session</code> and <code>transaction</code> .
<code>RESERVE_POOL_SIZE</code>	0	How many additional connections to allow per pool. 0 disables.
<code>RESERVE_POOL_TIMEOUT</code>	5	If a client has not been serviced in this many seconds, pgbouncer ends the connection.
<code>QUERY_TIMEOUT</code>	0	Queries running longer than that are canceled.
<code>IGNORE_STARTUP_PARAMETERS</code>	<code>extra_float_digits</code>	Set to ignore particular parameters in startup packets.
<code>PG_PORT</code>	5432	The port to use when connecting to the database.
<code>CRUNCHY_DEBUG</code>	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal sensitive information.

The crunchy-pgdump container executes either a `pg_dump` or `pg_dumpall` database backup against another PostgreSQL database.

Packages

The crunchy-pgdump Docker image contains the following packages (versions vary depending on PostgreSQL version):

- PostgreSQL (11.1, 10.6, 9.6.11 and 9.5.15)
- CentOS7 - publicly available
- RHEL7 - customers only

Environment Variables

Required

Name	Default	Description
PGDUMP_DB	None	Name of the database the backup is being performed on.
PGDUMP_HOST	None	Hostname of the database the backup is being performed on.
PGDUMP_PASS	None	Password for the PostgreSQL role being used.
PGDUMP_USER	None	Username for the PostgreSQL role being used.

Optional

Name	Default	Description
PGDUMP_ALL	TRUE	Run <code>pg_dump</code> instead of <code>pg_dumpall</code> . Set to false to enable <code>pg_dump</code> .
PGDUMP_CUSTOM_OPTS	None	Advanced options to pass into <code>pg_dump</code> or <code>pg_dumpall</code> .
PGDUMP_FILENAME	dump	Name of the file created by the <code>pgdump</code> container.
PGDUMP_PORT	5432	Port of the PostgreSQL database to connect to.
CRUNCHY_DEBUG	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal secrets in logs.

For a list of advanced options for configuring the `PGDUMP_CUSTOM_OPTS` variable, see the official documentation:

<https://www.postgresql.org/docs/current/static/app-pgdump.html>

<https://www.postgresql.org/docs/current/static/app-pg-dumpall.html>

Volumes

Name	Description
/pgdata	Volume used to store the data directory contents for the PostgreSQL database.

Dump Location

Backups are stored in a mounted backup volume location, using the database host name plus *-backups* as a sub-directory, then followed by a unique backup directory based upon a date/timestamp. It is left to the user to perform database backup archives in this current version of the container. This backup location is referenced when performing a database restore.

The crunchy-pgpool container executes the [pgPool II](#) utility. `pgPool` can be used to provide a smart PostgreSQL-aware proxy to a PostgreSQL cluster, both primary and replica, so that applications only have to work with a single database connection.

PostgreSQL replicas are read-only whereas a primary is capable of receiving both read and write actions.

The default `pgPool` examples use a Secret to hold the set of `pgPool` configuration files used by the examples. The Secret is mounted into the `pgconf` volume mount where the container will look to find configuration files. If you do not specify your own configuration files via a Secret then you can specify environment variables to the container that it will attempt to use to configure `pgPool`, although this is not recommended for production environments.

Features

The following features are supported by the `crunchy-postgres` container:

- Basic invocation of pgPool II

Packages

The `crunchy-pgpool` Docker image contains the following packages (versions vary depending on PostgreSQL version):

- PostgreSQL (11.1, 10.6, 9.6.11 and 9.5.15)
- [pgPool II](#)
- CentOS7 - publicly available
- RHEL7 - customers only

Environment Variables

Required

Name	Default	Description
<code>PG_USERNAME</code>	None	Username for the PostgreSQL role being used.
<code>PG_PASSWORD</code>	None	Password for the PostgreSQL role being used.
<code>PG_PRIMARY_SERVICE_NAME</code>	None	Database host to connect to for the primary node.
<code>PG_REPLICA_SERVICE_NAME</code>	None	Database host to connect to for the replica node.

Optional

Name	Default	Description
<code>CRUNCHY_DEBUG</code>	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal secrets in logs.

The `restore` image provides a means of performing a restore of a dump from `pg_dump` or `pg_dumpall` via `psql` or `pg_restore` to a PostgreSQL container database.

Packages

The `crunchy-pgrestore` Docker image contains the following packages (versions vary depending on PostgreSQL version):

- PostgreSQL (11.1, 10.6, 9.6.11 and 9.5.15)
- CentOS7 - publicly available
- RHEL7 - customers only

Environment Variables

Required

Name	Default	Description
<code>PGRESTORE_DB</code>	None	Name of the database to connect to.
<code>PGRESTORE_HOST</code>	None	Hostname of the database to connect to.
<code>PGRESTORE_PASS</code>	None	Password for the PostgreSQL role being used.
<code>PGRESTORE_USER</code>	None	Username for the PostgreSQL role being used.

Optional

Name	Default	Description
<code>PGDUMP_BACKUP_HOST</code>	None	Hostname of the PostgreSQL server that was backed up by <code>pgdump</code> container.
<code>PGRESTORE_BACKUP_TIMESTAMP</code>	Empty	Timestamp of the backup to restore from.
<code>PGRESTORE_CUSTOM_OPTS</code>	Empty	Advanced options to pass into <code>pg_restore</code> .
<code>PGRESTORE_PORT</code>	5432	Port of the PostgreSQL database to connect to.
<code>CRUNCHY_DEBUG</code>	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal sensitive information.

For a list of advanced options for configuring the `PGRESTORE_CUSTOM_OPTS` variable, see the [official documentation](#).

Volumes

Name	Description
<code>/pgdata</code>	Volume used to store the data directory contents for the PostgreSQL database.

Dump-file Input Location

As the input files for `crunchy-pgrestore`, files generated by `crunchy-pgdump` are retrieved in a mounted backup volume location, using the database host name plus `-backups` as a sub-directory, then followed by a unique backup directory based upon a date/timestamp. It is left to the user to restore database dump archives in this current version of the container.

PostgreSQL (pronounced “post-gress-Q-L”) is an open source, ACID compliant, relational database management system (RDBMS) developed by a worldwide team of volunteers. The `crunchy-postgres-gis` container image is unmodified, open source PostgreSQL packaged and maintained by professionals. This image is identical to the `crunchy-postgres` image except it includes the open source geospatial extension [PostGIS](#) for PostgreSQL in addition to the language extension [PL/R](#) which allows for writing functions in the R statistical computing language.

Features

The following features are supported by the `crunchy-postgres-gis` container:

- Kubernetes and OpenShift secrets
- Backup and restoration from various tools: `pgbackrest`, `pg_basebackup` and `pg_dump/pg_restore`.
- Custom mounted configuration files (see below)
- Async and Sync Replication
- PostGIS
- PL/R

Packages

The `crunchy-postgres-gis` Docker image contains the following packages (versions vary depending on PostgreSQL version):

- PostgreSQL (11.1, 10.6, 9.6.11 and 9.5.15)
- [pgBackRest](#) (2.x)
- CentOS7 - publicly available
- RHEL7 - customers only

Environment Variables

Required

Name	Default	Description
PG_DATABASE	None	Set this value to create an initial database
PG_PRIMARY_PORT	None	Set this value to configure the primary PostgreSQL port. It is recommended to use 5432
PG_MODE	None	Set to primary , replica or set to specify the mode of the database
PG_USER	None	Set this value to specify the username of the general user account
PG_PASSWORD	None	Set this value to specify the password of the user role
PG_PRIMARY_USER	None	Set this value to specify the username of the replication user
PG_PRIMARY_PASSWORD	None	Set this value to specify the password of the replication user
PG_ROOT_PASSWORD	None	Set this value to specify the password of the superuser role

Optional

Name	Default	Description
ARCHIVE_MODE	Off	Set this value to on to enable continuous WAL archiving
ARCHIVE_TIMEOUT	60	Set to a number (in seconds) to configure archive_timeout in postgresql.conf
CHECKSUMS	Off	Enables data-checksums during initialization of the database. Can only be used with PostgreSQL 12 and later.
CRUNCHY_DEBUG	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal sensitive information.
LOG_STATEMENT	none	Sets the log_statement value in postgresql.conf
LOG_MIN_DURATION_STATEMENT	60000	Sets the log_min_duration_statement value in postgresql.conf
MAX_CONNECTIONS	100	Sets the max_connections value in postgresql.conf
MAX_WAL_SENDERS	6	Set this value to configure the max number of WAL senders (replication)
PG_LOCALE	UTF-8	Set the locale of the database
PG_PRIMARY_HOST	None	Set this value to specify primary host. Note: only used when PG_MODE != primary
PG_REPLICA_HOST	None	Set this value to specify the replica host label. Note; used when PG_MODE is replica
PGAUDIT_ANALYZE	None	Set this to enable pgaudit_analyze
PGBOUNCER_PASSWORD	None	Set this to enable pgBouncer support by creating a special pgbouncer user
PGDATA_PATH_OVERRIDE	None	Set this value to override the /pgdata directory name. By default /pgdata
SHARED_BUFFERS	128MB	Set this value to configure shared_buffers in postgresql.conf
SYNC_REPLICA	None	Set this value to specify the names of replicas that should use synchronized replication
TEMP_BUFFERS	8MB	Set this value to configure temp_buffers in postgresql.conf
WORK_MEM	4MB	Set this value to configure work_mem in postgresql.conf
XLOGDIR	None	Set this value to configure PostgreSQL to send WAL to the /pgwal volume
PGBACKREST	false	Set this value to true in order to enable and initialize pgBackRest in the container
BACKREST_SKIP_CREATE_STANZA	false	Set this value to true in order to skip the configuration check and the automatic creation of the backrest stanza

Volumes

Name	Description
/backrestrepo	Volume used by the pgbackrest backup tool to store physical backups.
/backup	Volume used by the pg_basebackup backup tool to store physical backups.
/pgconf	Volume used to store custom configuration files mounted to the container.
/pgdata	Volume used to store the data directory contents for the PostgreSQL database.
/pgwal	Volume used to store Write Ahead Log (WAL) when XLOGDIR environment variable is set to true .
/recover	Volume used for Point In Time Recovery (PITR) during startup of the PostgreSQL database.

Custom Configuration

The following configuration files can be mounted to the `/pgconf` volume in the `crunchy-postgres` container to customize the runtime:

Name	Description
<code>ca.crt</code>	Certificate of the CA used by the server when using SSL authentication
<code>ca.crl</code>	Revocation list of the CA used by the server when using SSL authentication
<code>pg_hba.conf</code>	Client authentication rules for the database
<code>pg_ident.conf</code>	Mapping of external users (such as SSL certs, GSSAPI, LDAP) to database users
<code>postgresql.conf</code>	PostgreSQL settings
<code>server.key</code>	Key used by the server when using SSL authentication
<code>server.crt</code>	Certificate used by the server when using SSL authentication
<code>setup.sql</code>	Custom SQL to execute against the database. Note: only run during the first startup (initialization)

Verifying PL/R

In order to verify the successful initialization of the PL/R extension, the following commands can be run:

```
create extension plr;
SELECT * FROM plr_environ();
SELECT load_r_typenames();
SELECT * FROM r_typenames();
SELECT plr_array_accum('{23,35}', 42);
CREATE OR REPLACE FUNCTION plr_array (text, text)
RETURNS text []
AS '$libdir/plr','plr_array'
LANGUAGE 'c' WITH (isstrict);
select plr_array('hello','world');
```

PostgreSQL (pronounced “post-gress-Q-L”) is an open source, ACID compliant, relational database management system (RDBMS) developed by a worldwide team of volunteers. The `crunchy-postgres` container image is unmodified, open source PostgreSQL packaged and maintained by professionals.

Features

The following features are supported by the `crunchy-postgres` container:

- Kubernetes and OpenShift secrets
- Backup and restoration from various tools: `pgbackrest`, `pg_basebackup` and `pg_dump/pg_restore`.
- Custom mounted configuration files (see below)
- Async and Sync Replication

Packages

The `crunchy-postgres` Docker image contains the following packages (versions vary depending on PostgreSQL version):

- PostgreSQL (11.1, 10.6, 9.6.11 and 9.5.15)
- [pgBackRest](#) (2.x)
- CentOS7 - publicly available
- RHEL7 - customers only

Environment Variables

Required

Name	Default	Description
PG_DATABASE	None	Set this value to create an initial database
PG_PRIMARY_PORT	None	Set this value to configure the primary PostgreSQL port. It is recommended to use 5432
PG_MODE	None	Set to primary , replica or set to specify the mode of the database
PG_USER	None	Set this value to specify the username of the general user account
PG_PASSWORD	None	Set this value to specify the password of the user role
PG_PRIMARY_USER	None	Set this value to specify the username of the replication user
PG_PRIMARY_PASSWORD	None	Set this value to specify the password of the replication user
PG_ROOT_PASSWORD	None	Set this value to specify the password of the superuser role

Optional

Name	Default	Description
ARCHIVE_MODE	Off	Set this value to on to enable continuous WAL archiving
ARCHIVE_TIMEOUT	60	Set to a number (in seconds) to configure archive_timeout in postgresql.conf
CHECKSUMS	true	Enables data-checksums during initialization of the database. Can only be set to off if archive_mode is on
CRUNCHY_DEBUG	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal sensitive information
LOG_STATEMENT	none	Sets the log_statement value in postgresql.conf
LOG_MIN_DURATION_STATEMENT	60000	Sets the log_min_duration_statement value in postgresql.conf
MAX_CONNECTIONS	100	Sets the max_connections value in postgresql.conf
MAX_WAL_SENDERS	6	Set this value to configure the max number of WAL senders (replication)
PG_LOCALE	UTF-8	Set the locale of the database
PG_PRIMARY_HOST	None	Set this value to specify primary host. Note: only used when PG_MODE != primary
PG_REPLICA_HOST	None	Set this value to specify the replica host label. Note; used when PG_MODE is replica
PGAUDIT_ANALYZE	None	Set this to enable pgaudit_analyze
PGBOUNCER_PASSWORD	None	Set this to enable pgBouncer support by creating a special pgbouncer user
PGMONITOR_PASSWORD	None	Set this to enable pgMonitor support by creating a special ccp_monitoring user
PGDATA_PATH_OVERRIDE	None	Set this value to override the /pgdata directory name. By default /pgdata
SHARED_BUFFERS	128MB	Set this value to configure shared_buffers in postgresql.conf
SYNC_REPLICA	None	Set this value to specify the names of replicas that should use synchronized replication
TEMP_BUFFERS	8MB	Set this value to configure temp_buffers in postgresql.conf
WORK_MEM	4MB	Set this value to configure work_mem in postgresql.conf
XLOGDIR	None	Set this value to configure PostgreSQL to send WAL to the /pgwal volume
PGBACKREST	false	Set this value to true in order to enable and initialize pgBackRest in the container
BACKREST_SKIP_CREATE_STANZA	false	Set this value to true in order to skip the configuration check and the automatic creation of the backrest stanza

Volumes

Name	Description
/backrestrepo	Volume used by the pgbackrest backup tool to store physical backups.
/backup	Volume used by the pg_basebackup backup tool to store physical backups.
/pgconf	Volume used to store custom configuration files mounted to the container.
/pgdata	Volume used to store the data directory contents for the PostgreSQL database.
/pgwal	Volume used to store Write Ahead Log (WAL) when XLOGDIR environment variable is set to true .

Name	Description
<code>/recover</code>	Volume used for Point In Time Recovery (PITR) during startup of the PostgreSQL database.

Custom Configuration

The following configuration files can be mounted to the `/pgconf` volume in the `crunchy-postgres` container to customize the runtime:

Name	Description
<code>ca.crt</code>	Certificate of the CA used by the server when using SSL authentication
<code>ca.crl</code>	Revocation list of the CA used by the server when using SSL authentication
<code>pg_hba.conf</code>	Client authentication rules for the database
<code>pg_ident.conf</code>	Mapping of external users (such as SSL certs, GSSAPI, LDAP) to database users
<code>postgresql.conf</code>	PostgreSQL settings
<code>server.key</code>	Key used by the server when using SSL authentication
<code>server.crt</code>	Certificate used by the server when using SSL authentication
<code>setup.sql</code>	Custom SQL to execute against the database. Note: only run during the first startup (initialization)

[Prometheus](#) is a multi-dimensional time series data model with an elastic query language. It is used in collaboration with [Grafana](#) in this metrics suite. Overall, it's reliable, manageable, and operationally simple for efficiently storing and analyzing data for large-scale environments. It scrapes metrics from exporters such as the ones utilized by the `crunchy-collect` container. The `crunchy-prometheus` container must be able to reach the `crunchy-collect` container in order to scrape metrics.

By default, `crunchy-prometheus` detects which environment its running on (Docker, Kubernetes, or OpenShift) and applies a default configuration. If this container is running on Kubernetes or OpenShift, it will use the Kubernetes API to discover pods with the label `"crunchy-collect": "true"`. The `crunchy-collect` container must have this label defined in order to be discovered.

For Docker environments the `crunchy-collect` hostname must be specified as an environment variable.

A user may define a custom `prometheus.yml` file and mount to `/conf` for custom configuration. For configuration examples, see [here](#).

The following port is exposed by the `crunchy-prometheus` container:

- `crunchy-prometheus:9090` - the Prometheus web user interface

Packages

The `crunchy-prometheus` Docker image contains the following packages:

- [Prometheus](#)
- CentOS7 - publicly available
- RHEL7 - customers only

Environment Variables

Required

Name	Default	Description
<code>COLLECT_HOST</code>	None	Hostname of Crunchy Collect container. Only required in Docker environments.

Optional

Name	Default	Description
SCRAPE_INTERVAL	5s	Set this value to the number of seconds to scrape metrics from exporters.
SCRAPE_TIMEOUT	5s	Set this value to the number of seconds to timeout when scraping metrics from exporters.
CRUNCHY_DEBUG	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal secrets in logs.

Permissions

Crunchy Prometheus queries Kubernetes to discover metric exporters and additional Kubernetes metadata. Due to the integration with Kubernetes, Crunchy Prometheus requires a service account with the following permissions:

- Cluster Role
- Pods: `get`, `list`, `watch`

The Crunchy Scheduler container provides a cronlike microservice for automating `pgBaseBackup` and `pgBackRest` backups within a single namespace. The scheduler watches Kubernetes for config maps with the label `crunchy-scheduler=true`. If found the scheduler parses a JSON object contained in the config map and converts it into an scheduled task.

Packages

The Crunchy Scheduler Docker image contains the following packages:

- CentOS7 - publicly available
- RHEL7 - customers only
- Scheduler App

Environment Variables

Required

Name	Default	Description
NAMESPACE	None	The namespace the microservice should watch. Crunchy Scheduler only works in a single namespace.
TIMEOUT	300	The time (in seconds) the scheduler should wait before timing out on a backup job.

Optional

Name	Default	Description
CRUNCHY_DEBUG	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal secrets in logs.

Permissions

Crunchy Scheduler queries Kubernetes to discover schedules and perform scheduled tasks (either creating a job or running commands against a PostgreSQL container). Due to the integration with Kubernetes, Crunchy Scheduler requires a service account with the following permissions:

- Role
- ConfigMaps: `get`, `list`, `watch`
- Deployments: `get`, `list`, `watch`
- Jobs: `get`, `list`, `watch`, `create`, `delete`
- Pods: `get`, `list`, `watch`
- Pods/Exec: `create`
- Secrets: `get`, `list`, `watch`

Timezone

Crunchy Scheduler uses the UTC timezone for all schedules.

Schedule Expression Format

Schedules are expressed using the following rules:

Field name	Mandatory?	Allowed values	Allowed special characters
Seconds	Yes	0-59	* / , -
Minutes	Yes	0-59	* / , -
Hours	Yes	0-23	* / , -
Day of month	Yes	1-31	* / , - ?
Month	Yes	1-12 or JAN-DEC	* / , -
Day of week	Yes	0-6 or SUN-SAT	* / , - ?

The crunchy-upgrade container contains both the 9.5 / 9.6 and 9.6 / 10 PostgreSQL packages in order to perform a pg_upgrade from 9.5 to 9.6 or 9.6 to 10 versions.

Features

The following features are supported by the crunchy-upgrade container:

- Supports a pg_upgrade of the PostgreSQL database.
- Doesn't alter the old database files.
- Creates the new database directory.

Restrictions

- Does **not** currently support a PostGIS upgrade.
- Supports upgrades from only 9.5 to 9.6, or 9.6 to 10.

Packages

The crunchy-upgrade Docker image contains the following packages (versions vary depending on PostgreSQL version):

- PostgreSQL (11.1, 10.6, 9.6.11 and 9.5.15)
- CentOS7 - publicly available
- RHEL7 - customers only

Environment Variables

Required

Name	Default	Description
OLD_DATABASE_NAME	None	Refers to the database (pod) name that we want to convert.
NEW_DATABASE_NAME	None	Refers to the database (pod) name that is given to the upgraded database.
OLD_VERSION	None	The PostgreSQL version of the old database.
NEW_VERSION	None	The PostgreSQL version of the new database.

Optional

Name	Default	Description
PG_LOCALE	Default locale	If set, the locale you want to create the database with.

Name	Default	Description
CHECKSUMS	true	Enables <code>data-checksums</code> during initialization of the database. Can only be set during initial
XLOGDIR	None	If set, <code>initdb</code> will use the specified directory for WAL.
CRUNCHY_DEBUG	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal secrets in logs.

Data checksums on the Crunchy PostgreSQL container were enabled by default in version 2.1.0. When trying to upgrade, it's required that both the old database and the new database have the same data checksums setting. Prior to upgrade, check if `data_checksums` were enabled on the database by running the following SQL: `SHOW data_checksums`

The `crunchy-vacuum` container allows you to perform a SQL `VACUUM` job against a PostgreSQL database container. You specify a database to vacuum using various environment variables which are listed below. It is possible to run different vacuum operations either manually or automatically through scheduling.

The `crunchy-vacuum` image is executed with the PostgreSQL connection parameters passed to the single-primary PostgreSQL container. The type of vacuum performed is dictated by the environment variables passed into the job.

More information on the PostgreSQL `VACUUM` job can be found in the [official PostgreSQL documentation](#).

Packages

The `crunchy-vacuum` Docker image contains the following packages:

- CentOS7 - publicly available
- RHEL7 - customers only

Environment Variables

Required

Name	Default	Description
JOB_HOST	None	The PostgreSQL host the <code>VACUUM</code> should be performed against.
PG_DATABASE	None	The PostgreSQL database the <code>VACUUM</code> should be performed against.
PG_PORT	5432	Allows you to override the default value of 5432.
PG_USER	None	Username for the PostgreSQL role being used.
PG_PASSWORD	None	Password for the PostgreSQL role being used.

Optional

Name	Default	Description
VAC_FULL	TRUE	When set to true, adds the <code>FULL</code> parameter to the <code>VACUUM</code> command.
VAC_TABLE	FALSE	When set to true, allows you to specify a single table to vacuum. When not specified, the entire da
VAC_ANALYZE	TRUE	When set to true, adds the <code>ANALYZE</code> parameter to the <code>VACUUM</code> command.
VAC_VERBOSE	TRUE	When set to true, adds the <code>VERBOSE</code> parameter to the <code>VACUUM</code> command.
VAC_FREEZE	FALSE	When set to true, adds the <code>FREEZE</code> parameter to the <code>VACUUM</code> command.
CRUNCHY_DEBUG	FALSE	Set this to true to enable debugging in logs. Note: this mode can reveal secrets in logs.

pgAdmin4 example

This example deploys the `pgadmin4 v2` web user interface for PostgreSQL without TLS.

After running the example, you should be able to browse to `http://127.0.0.1:5050` and log into the web application with the following configured credentials:

- Username : *admin@admin.com*
- Password: *password*

If you are running this example using Kubernetes or OpenShift, it is required to use a port-forward proxy to access the dashboard.

To start the port-forward proxy run the following:

```
`${CCP_CLI}` port-forward pgadmin4-http 5050:5050
```

To access the pgAdmin4 dashboard through the proxy, navigate to *http://127.0.0.1:5050* in a browser.

See the link:<http://pgadmin.org>[pgAdmin4 documentation] for more details.

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

To run this example, run the following:

```
cd $CCPROOT/examples/docker/pgadmin4-http  
./run.sh
```

Kubernetes and OpenShift

Start the container as follows:

```
cd $CCPROOT/examples/kube/pgadmin4-http  
./run.sh
```

An emptyDir with write access must be mounted to the */run/httpd* directory in OpenShift.

pgAdmin4 with TLS

This example deploys the pgadmin4 v2 web user interface for PostgreSQL with TLS.

After running the example, you should be able to browse to *https://127.0.0.1:5050* and log into the web application with the following configured credentials:

- Username : *admin@admin.com*
- Password: *password*

If you are running this example using Kubernetes or OpenShift, it is required to use a port-forward proxy to access the dashboard.

To start the port-forward proxy run the following:

```
`${CCP_CLI}` port-forward pgadmin4-https 5050:5050
```

To access the pgAdmin4 dashboard through the proxy, navigate to *https://127.0.0.1:5050* in a browser.

See the link:<http://pgadmin.org>[pgadmin4 documentation] for more details.

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

To run this example, run the following:

```
cd $CCPROOT/examples/docker/pgadmin4-https  
./run.sh
```

Kubernetes and OpenShift

Start the container as follows:

```
cd $CCPROOT/examples/kube/pgadmin4-https
./run.sh
```

An emptyDir with write access must be mounted to the `/run/httpd` directory in OpenShift.

Major Upgrade

This example assumes you have run *primary* using a PG 9.5 or 9.6 image such as `centos7-9.5.15-2.2.0` prior to running this upgrade.

Starting in release 1.3.1, the upgrade container will let you perform a `pg_upgrade` either from a PostgreSQL version 9.5 database to 9.6 or from 9.6 to 10.

Prior to running this example, make sure your `CCP_IMAGE_TAG` environment variable is using the next major version of PostgreSQL that you want to upgrade to. For example, if you're upgrading from 9.5 to 9.6, make sure the variable references a PG 9.6 image such as `centos7-9.6.11-2.2.0`.

This will create the following in your Kubernetes environment:

- a Kubernetes Job running the *crunchy-upgrade* container
- a new data directory name *upgrade* found in the *pgnewdata* PVC

Data checksums on the Crunchy PostgreSQL container were enabled by default in version 2.1.0. When trying to upgrade, it's required that both the old database and the new database have the same data checksums setting. Prior to upgrade, check if `data_checksums` were enabled on the database by running the following SQL: `SHOW data_checksums`

Kubernetes and OpenShift

Before running the example, ensure you edit `upgrade.json` and update the `OLD_VERSION` and `NEW_VERSION` parameters to the ones relevant to your situation.

Start the upgrade as follows:

```
cd $CCPROOT/examples/kube/upgrade
./run.sh
```

If successful, the Job will end with a **successful** status. Verify the results of the Job by examining the Job's pod log:

```
${CCP_CLI} get pod -l job-name=upgrade
${CCP_CLI} logs -l job-name=upgrade
```

You can verify the upgraded database by running the `post-upgrade.sh` script in the `examples/kube/upgrade` directory. This will create a PostgreSQL pod that mounts the upgraded volume.

pgBackRest Examples

Written and maintained by David Steele, pgBackRest is a utility that provides backup and restore functionality for PostgreSQL databases. pgBackRest is available for use within the Crunchy Container Suite, and can therefore be utilized to provide an effective backup and restore solution for any database clusters deployed using the `crunchy-postgres` or `crunchy-postgres-gis` containers. The following section will provide an overview of how pgBackRest can be utilized within the Crunchy Container Suite, including examples for enabling and configuring pgBackRest, and then utilizing pgBackRest to backup and restore various PostgreSQL database clusters. For more detailed information about pgBackRest, please visit the [official pgBackRest website](#).

Configuration Overview

In order to enable pgBackRest within a `crunchy-postgres` or `crunchy-postgres-gis` container, environment variable `PGBACKREST` must be set to `true` during deployment of the container (`PGBACKREST=true`). This will setup the proper pgBackRest configuration, ensure any required pgBackRest repositories and directories are created, and will create the proper pgBackRest stanza.

Please note that setting `PGBACKREST=true` is all that is needed to configure and enable pgBackRest within a `crunchy-postgres` or `crunchy-postgres-gis` container. When enabled, default environment variables will be set for pgBackRest as follows, unless they are otherwise explicitly defined and provided during deployment of the container:

```
export PGBACKREST_STANZA="db"
export PGBACKREST_PG1_PATH="/pgdata/${PGDATA_DIR}"
export PGBACKREST_REPO1_PATH="/backrestrepo/${PGDATA_DIR}-backups"
export PGBACKREST_LOG_PATH="/tmp"
```

As shown above, a stanza named `db` is created by default, using the default values provided for both `PGBACKREST_PG1_PATH` and `PGBACKREST_REPO1_PATH`. Variable `PGDATA_DIR` represents the name of the database cluster's data directory, which will either be the hostname of the container or the value specified for variable `PGDATA_PATH_OVERRIDE` during deployment of the container. Please see the [crunchy-postgres](#) and/or [crunchy-postgres-gis](#) container specifications for additional details.

While setting `PGBACKREST` to `true` provides a simple method for enabling `pgBackRest` within a `crunchy-postgres` or `crunchy-postgres-gis` container, `pgBackRest` is also fully configurable and customizable via the various environment variables supported by `pgBackRest`. This applies to the `crunchy-backrest-restore` container as well, which is also configured using `pgBackRest` environment variables when performing database restores. Therefore, during the deployment of any container containing `pgBackRest` (`crunchy-postgres`, `crunchy-postgres-gis` or `crunchy-backrest-restore`), environment variables should be utilized to configure and customize the `pgBackRest` utility as needed and ensure the desired backup and restore functionality is achieved. For instance, the following environment variables could be specified upon deployment of the `crunchy-backrest-restore` container in order to perform delta restore to a specific point-in-time:

```
PGBACKREST_TYPE=time
PITR_TARGET="2018-12-27 16:53:05.590156+00"
PGBACKREST_DELTA=y
```

Full, incremental and differential backups of PostgreSQL databases deployed using the Crunchy Container Suite can be scheduled using `pgBackRest` and the `crunchy-scheduler` container, and/or can also be performed manually by executing `pgBackRest` commands against the desired `crunchy-postgres` or `crunchy-postgres-gis` container. Database restores, on the other hand, can be performed via the `crunchy-backrest-restore` container, which offers full `pgBackRest` restore capabilities, such as full, point-in-time and delta restores. Further information and guidance for performing both backups and restores using the Crunchy Container Suite and `pgBackRest` will be provided in the examples below. Additionally, for more information on utilizing the `crunchy-scheduler` container to schedule and perform `pgBackRest` database backups, please see the `crunchy-scheduler` [specifications](#) and [examples](#).

In addition to providing the backup and restoration capabilities discussed above, `pgBackRest` supports the capability to asynchronously push and get write ahead logs (WAL) to and from a WAL archive. To enable asynchronous WAL archiving within a `crunchy-postgres` or `crunchy-postgres-gis` container, `pgBackRest` environment variable `PGBACKREST_ARCHIVE_ASYNC` must be set to `"y"` during deployment (`PGBACKREST_ARCHIVE_ASYNC=y`). This will automatically enable WAL archiving within the container if not otherwise explicitly enabled, set the proper `pgbackrest archive` command within the `postgresql.conf` configuration file, and ensure the proper spool path has been created.

If a spool path is not explicitly provided using environment variable `PGBACKREST_SPOOL_PATH`, this variable will default as follows:

```
# Environment variable XLOGDIR=true"
export PGBACKREST_SPOOL_PATH="/pgdata/${PGDATA_DIR}"

# Environment variable XLOGDIR!=true
export PGBACKREST_SPOOL_PATH="/pgwal/${PGDATA_DIR}/spool"
```

As shown above, the default location of the spool path depends on whether or not `XLOGDIR=true`, with `XLOGDIR` enabling the storage of WAL to the `/pgwal` volume within the container. Being that `pgBackRest` recommends selecting a spool path that is as close to the WAL as possible, this provides a sensible default for the spool directory. However, `PGBACKREST_SPOOL_PATH` can also be explicitly configured during deployment to any path desired. And once again, `PGDATA_DIR` represents either the hostname of the container or the value specified for variable `PGDATA_PATH_OVERRIDE`.

The examples below will demonstrate the `pgBackRest` backup, restore and asynchronous archiving capabilities described above, while also providing insight into the proper configuration of `pgBackRest` within the Crunchy Container Suite. For more information on these `pgBackRest` capabilities and associated configuration, please consult the [official pgBackRest documentation](#).

Kubernetes and OpenShift

The `pgBackRest` examples for Kubernetes and OpenShift can be configured to use the PostGIS images by setting the following environment variable when running the examples:

```
export CCP_PG_IMAGE='-gis'
```

Backup

In order to demonstrate the backup and restore capabilities provided by `pgBackRest`, it is first necessary to deploy a PostgreSQL database, and then create a full backup of that database. This example will therefore deploy a `crunchy-postgres` or `crunchy-postgres-gis` container containing a PostgreSQL database, which will then be backed up manually by executing a `pgbackrest backup` command. **Please note**

that this example serves as a prerequisite for the restore examples that follow, and therefore must be run prior to running those examples.

Start the example as follows:

```
cd $CCPROOT/examples/kube/backrest/backup
./run.sh
```

This will create the following in your Kubernetes environment: - A pod named **backrest** containing a PostgreSQL database with pgBackRest configured - A service named **backrest** for the PostgreSQL database - A PV and PVC for the PGDATA directory - A PV and PVC for the pgBackRest backups and archives directories

Once the **backrest** pod is running, use the `pgbackrest info` command to verify that pgbackrest has been properly configured and WAL archiving is working properly:

```
$ ${CCP_CLI} exec backrest -- pgbackrest info \
  --stanza=db \
  --repo1-path=/backrestrepo/backrest-backups

pg_pid=126
stanza: db
  status: error (no valid backups)
  cipher: none

  db (current)
    wal archive min/max (11-1): 0000000100000000000000001 / 0000000100000000000000003
```

An output similar to the above indicates that pgBackRest was properly configured upon deployment of the pod, the **db** stanza has been created, and WAL archiving is working properly. The error next to **status** is expected being that a backup has not yet been generated.

Now that we have verified that pgBackRest is properly configured and enabled, a backup of the database can be generated. Being that this is the first backup of the database, we will take create a **full** backup:

```
$ ${CCP_CLI} exec backrest -- pgbackrest backup \
  --stanza=db \
  --pg1-path=/pgdata/backrest \
  --repo1-path=/backrestrepo/backrest-backups \
  --log-path=/tmp \
  --type=full

pg_pid=138
WARN: option repo1-retention-full is not set, the repository may run out of space
      HINT: to retain full backups indefinitely (without warning), set option
            'repo1-retention-full' to the maximum.
```

The warning displayed is expected, since backup retention has not been configured for this example. Assuming no errors are displayed, a full backup has now been successfully created.

Restore

pgBackRest provides numerous methods and strategies for restoring a PostgreSQL database. The following section will demonstrate three forms of database restores that can be accomplished when using pgBackRest with the Crunchy Container Suite: - **Full**: restore all database files into an empty PGDATA directory - **point-in-time Recovery (PITR)**: restore a database to a specific point-in-time using an empty PGDATA directory - **Delta**: restore a database to a specific point-in-time using an existing PGDATA directory

Full This example will demonstrate a full database restore to an empty PGDATA directory. *Please ensure the Backup example is currently running and a full backup has been generated prior to running this example.*

Prior to running the full restore, we will first make a change to the currently running database, which will we will then verify still exists following the restore. Create a simple table in the database as follows:

```
$ ${CCP_CLI} exec backrest -- psql -c "create table backrest_test_table (id int)"
CREATE TABLE
```

Now verify that the new table exists:

```
$ ${CCP_CLI} exec backrest -- psql -c "table backrest_test_table"
 id
----
(0 rows)
```

With the table in place, we can now start the full restore as follows:

```
cd $CCPROOT/examples/kube/backrest/full
./run.sh
```

This will create the following in your Kubernetes environment: - A Kubernetes job named **backrest-full-restore-job** which will perform the restore using the crunchy-backrest-restore container - A PV and PVC for the new PGDATA directory that will contain the restored database. The directory will initially be empty, as required pgBackRest when performing a full restore, and will then contain the restored database upon completion of the restore.

Please note that a brand new PV and PVC are created when running the restore to clearly indicate that the database will be restored into an entirely new (i.e. empty) volume as required by pgBackRest. The names of the new PV and PVC are as follows: - **PV:** `${CCP_NAMESPACE}-br-new-pgdata` - **PVC:** `br-new-pgdata`

You can verify that the restore has completed successfully by verifying that the Kubernetes job has completed successfully:

```
$ ${CCP_CLI} get jobs
NAME                                COMPLETIONS  DURATION  AGE
backrest-full-restore-job          1/1           15s       58s
```

Once the job is complete, the post restore script can then be run, which will create a new pod named **backrest-full-restored** containing the restored database:

```
cd $CCPROOT/examples/kube/backrest/full
./post-restore.sh
```

Finally, once the **backrest-full-restored** pod is running we can verify that the restore was successful by verifying that the table created prior to the restore still exists:

```
$ ${CCP_CLI} exec backrest-full-restored -- psql -c "table backrest_test_table"
 id
----
(0 rows)
```

Please note that the default behavior of pgBackRest is to recover to the end of the WAL archive stream, which is why the full restore contained all changes made since the initial full backup was taken, including the creation of table **backrest_test_table**. pgBackRest therefore played the entire WAL archive stream for all changes that occurred up until the restore.

*As a reminder, please remember to run the cleanup script for the **Backup** example after running the cleanup script for this example.*

PITR As demonstrated with the full restore above, the default behavior of pgBackRest is to recover to the end of the WAL archive stream. However, pgBackRest also provides the ability to recover to a specific point-in-time utilizing the WAL archives created since the last backup. This example will demonstrate how pgBackRest can be utilized to perform a point-in-time recovery (PITR) and therefore recover the database to specific point-in-time specified by the user. *Please ensure that the Backup example is currently running and a full backup has been generated prior to running this example.*

Prior to running the PITR restore, we will first verify the current state of the database, after which we will then make a change to the database. This will allow us to verify that the PITR is successful by providing a method of verifying that the database has been restored to its current state following the restore.

To verify the current state of the database, we will first verify that a table called **backrest_test_table** does not exist in the database.

```
$ ${CCP_CLI} exec backrest -- psql -c " table backrest_test_table"
ERROR:  relation "backrest_test_table" does not exist
LINE 1:  table backrest_test_table
          ^
command terminated with exit code 1
```

Next, capture the current timestamp, which will be used later in the example when performing the restore:

```
$ ${CCP_CLI} exec backrest -- psql -c "select current_timestamp"
 current_timestamp
-----
2018-12-27 16:53:05.590156+00
(1 row)
```

Now create table **backrest_test_table**:

```
$ ${CCP_CLI} exec backrest -- psql -c "create table backrest_test_table (id int)"
CREATE TABLE
```

Then verify that the new table exists:

```
$ ${CCP_CLI} exec backrest -- psql -c "table backrest_test_table"
id
----
(0 rows)
```

With the table in place, we can now start the PITR. However, the timestamp captured above must also be provided in order to instruct pgBackRest to recover to that specific point-in-time. This is done using the `CCP_BACKREST_TIMESTAMP` variable, which allows us to then start the PITR as follows (replace the timestamp in the command below with the timestamp you captured above):

```
cd $CCPROOT/examples/kube/backrest/pitr
CCP_BACKREST_TIMESTAMP="2018-12-20 09:49:02.275701+00" ./run.sh
```

This will create the following in your Kubernetes environment: - A Kubernetes job named **backrest-pitr-restore-job** which will perform the restore using the `crunchy-backrest-restore` container

Additionally, when this example is run, the following pgBackRest environment variables are provided to the `crunchy-backrest-restore` container in order to initiate PITR restore to the point-in-time specified by the timestamp (in addition to any other pgBackRest variables required by the Crunchy Container Suite and pgBackRest):

```
PGBACKREST_TYPE=time
PITR_TARGET="${CCP_BACKREST_TIMESTAMP}"
```

As can be seen above, the timestamp provided for `CCP_BACKREST_TIMESTAMP` is used to populate variable `PITR_TARGET`, and therefore specify the point-in-time to restore the database to, while `PGBACKREST_TYPE` is set to `time` to indicate that a PITR should be performed.

Please note that the following pgBackRest environment variable is also set when performing the PITR, which results in a restore to a new/empty directory within an existing PV:

```
PGBACKREST_PG1_PATH=/pgdata/backrest-pitr-restored
```

You can verify that the restore has completed successfully by verifying that the Kubernetes job has completed successfully:

```
$ ${CCP_CLI} get jobs
NAME                                COMPLETIONS  DURATION  AGE
backrest-pitr-restore-job           1/1           15s       58s
```

Once the job is complete, the post restore script can then be run, which will create a new pod named **backrest-pitr-restored** containing the restored database:

```
cd $CCPROOT/examples/kube/backrest/pitr
./post-restore.sh
```

Finally, once the **backrest-pitr-restored** pod is running we can verify that the restore was successful by verifying that the table created prior to the restore no longer exists:

```
$ ${CCP_CLI} exec backrest-pitr-restored -- psql -c " table backrest_test_table"
ERROR:  relation "backrest_test_table" does not exist
LINE 1:  table backrest_test_table
          ^
command terminated with exit code 1
```

*As a reminder, please remember to run the cleanup script for the **Backup** example after running the cleanup script for this example.*

Delta By default, pgBackRest requires a clean/empty directory in order to perform a restore. However, pgBackRest also provides an another option when performing the restore in the form of the **delta** option, which allows the restore to be run against an existing PGDATA directory. With the delta option enabled, pgBackRest will use checksums to determine which files in the directory can be preserved, and which need to be restored (please note that pgBackRest will also remove any files that are not present in the backup). This example will again demonstrate a point-in-time recovery (PITR), only this time the restore will occur within the existing PGDATA directory by specifying the **delta** option during the restore. *Please ensure that the Backup example is currently running and a full backup has been generated prior to running this example.*

Prior to running the delta restore, we will first verify the current state of the database, and we will then make a change to the database. This will allow us to verify that the delta restore is successful by providing a method of verifying that the database has been restored to its current state following the restore.

To verify the current state of the database, we will first verify that a table called **backrest_test_table** does not exist in the database.

```
$ ${CCP_CLI} exec backrest -- psql -c " table backrest_test_table"
ERROR:  relation "backrest_test_table" does not exist
LINE 1:  table backrest_test_table
           ^
command terminated with exit code 1
```

Next, capture the current timestamp, which will be used later in the example when performing the restore:

```
$ ${CCP_CLI} exec backrest -- psql -c "select current_timestamp"
current_timestamp
-----
2018-12-27 16:53:05.590156+00
(1 row)
```

Now create table **backrest_test_table**:

```
$ ${CCP_CLI} exec backrest -- psql -c "create table backrest_test_table (id int)"
CREATE TABLE
```

Then verify that the new table exists:

```
$ ${CCP_CLI} exec backrest -- psql -c "table backrest_test_table"
 id
----
(0 rows)
```

With the table in place, we can now start the delta restore. When running the restore example the timestamp captured above must also be provided in order to instruct pgBackRest to recover to that specific point-in-time. This is done using the **CCP_BACKREST_TIMESTAMP** variable, which allows us to then start the delta restore as follows (replace the timestamp in the command below with the timestamp you captured above):

```
cd $CCPROOT/examples/kube/backrest/delta
CCP_BACKREST_TIMESTAMP="2018-12-20 09:49:02.275701+00" ./run.sh
```

This will create the following in your Kubernetes environment: - A Kubernetes job named **backrest-delta-restore-job** which will perform the restore using the **crunchy-backrest-restore** container

Additionally, when this example is run, the following pgBackRest environment variables are provided to the **crunchy-backrest-restore** container in order to initiate a delta restore to the point-in-time specified by the timestamp (in addition to any other pgBackRest variables required by the Crunchy Container Suite and pgBackRest):

```
PGBACKREST_TYPE=time
PITR_TARGET="${CCP_BACKREST_TIMESTAMP}"
PGBACKREST_DELTA=y
```

As can be seen above, the timestamp provided for **CCP_BACKREST_TIMESTAMP** is used to populate variable **PITR_TARGET**, and therefore specify the point-in-time to restore to, while **PGBACKREST_TYPE** is set to **time** to indicate that a PITR should be performed. **PGBACKREST_DELTA** is set to **y** to indicate that the delta option should be utilized when performing the restore.

It's also worth noting that the following pgBackRest environment variable is also set when performing the delta restore, which results in a restore within the existing **PGDATA** directory utilized by the database deployed when running the **Backup** example:

```
PGBACKREST_PG1_PATH=/pgdata/backrest
```

You can then verify that the restore has completed successfully by verifying that the Kubernetes job has completed successfully:

```
$ ${CCP_CLI} get jobs
NAME                                COMPLETIONS  DURATION  AGE
backrest-delta-restore-job          1/1           15s       58s
```

Once the job is complete, the post restore script can then be run, which will create a new pod named **backrest-delta-restored** containing the restored database:

```
cd $CCPROOT/examples/kube/backrest/delta
./post-restore.sh
```

Finally, once the **backrest-delta-restored** pod is running we can verify that the restore was successful by verifying that the table created prior to the restore no longer exists:

```
$ ${CCP_CLI} exec backrest-delta-restored -- psql -c " table backrest_test_table"
ERROR:  relation "backrest_test_table" does not exist
LINE 1:  table backrest_test_table
          ^
command terminated with exit code 1
```

As a reminder, please remember to run the cleanup script for the **Backup** example after running the cleanup script for this example.

Async Archiving

pgBackRest supports the capability to asynchronously push and get write ahead logs (WAL) to and from a WAL archive. Asynchronous archiving can improve performance by parallelizing operations, while also reducing the number of connections to remote storage. For more information on async archiving and its benefits, please see the [official pgBackRest documentation](#). This example will demonstrate how asynchronous archiving can be enabled within a crunchy-postgres or crunchy-postgres-gis container, while then also demonstrating the creation of a differential backup.

Start the example as follows:

```
cd $CCPROOT/examples/kube/backrest/async-archiving
./run.sh
```

This will create the following in your Kubernetes environment: - A pod named **backrest-async-archive** containing a PostgreSQL database with pgBackRest configured - A service named **backrest-async-archive** for the PostgreSQL database - A PV and PVC for the PGDATA directory - A PV and PVC for the pgBackRest backups and archives directories

Additionally, the following variable will be set during deployment of the pod in order to enable asynchronous archiving:

```
PGBACKREST_ARCHIVE_ASYNC=y
```

This will also result in the creation of the required pool path, which we can see by listing the contents of the `/pgdata` directory in the `backrest-async-archive` pod:

```
$ ${CCP_CLI} exec backrest-async-archive -- ls /pgdata
backrest-async-archive
backrest-async-archive-backups
backrest-async-archive-spool
```

Once the database is up and running, a full backup can be taken:

```
${CCP_CLI} exec backrest-async-archive -- pgbackrest backup \
--stanza=db \
--pg1-path=/pgdata/backrest-async-archive \
--repo1-path=/backrestrepo/backrest-async-archive-backups \
--log-path=/tmp \
--type=full
```

And once a full backup has been taken, other types of backups can also be taken using pgBackRest, such as a differential backup:

```
${CCP_CLI} exec backrest-async-archive -- pgbackrest backup \
--stanza=db \
--pg1-path=/pgdata/backrest-async-archive \
--repo1-path=/backrestrepo/backrest-async-archive-backups \
--log-path=/tmp \
--type=diff
```

The following command can then be run to verify that both backups were created successfully:

```
${CCP_CLI} exec backrest-async-archive -- pgbackrest info \
--stanza=db \
--repo1-path=/backrestrepo/backrest-async-archive-backups
```

Docker

Backup

In order to demonstrate the backup and restore capabilities provided by pgBackRest, it is first necessary to deploy a PostgreSQL database, and then create a full backup of that database. This example will therefore deploy a crunchy-postgres or crunchy-postgres-gis container

containing a PostgreSQL database, which will then be backed up manually by executing a `pgbackrest` backup command. *Please note that this example serves as a prerequisite for the restore examples that follow, and therefore must be run prior to running those examples.*

Start the example as follows:

```
cd $CCPROOT/examples/docker/backrest/backup
./run.sh
```

This will create the following in your Docker environment: - A container named **backrest** containing a PostgreSQL database with pgBackRest configured - A volume for the PGDATA directory - A volume for the pgBackRest backups and archives directories

Once the **backrest** container is running, use the `pgbackrest info` command to verify that pgbackrest has been properly configured and WAL archiving is working properly:

```
$ docker exec backrest pgbackrest info \
--stanza=db \
--repo1-path=/backrestrepo/backrest-backups

pg_pid=126
stanza: db
  status: error (no valid backups)
  cipher: none

  db (current)
    wal archive min/max (11-1): 00000001000000000000000001 / 00000001000000000000000003
```

An output similar to the above indicates that pgBackRest was properly configured upon deployment of the container, the **db** stanza has been created, and WAL archiving is working properly. The error next to **status** is expected being that a backup has not yet been generated.

Now that we have verified that pgBackRest is properly configured and enabled, a backup of the database can be generated. Being that this is the first backup of the database, we will take create a **full** backup:

```
$ docker exec backrest pgbackrest backup \
--stanza=db \
--pg1-path=/pgdata/backrest \
--repo1-path=/backrestrepo/backrest-backups \
--log-path=/tmp \
--type=full

pg_pid=138
WARN: option repo1-retention-full is not set, the repository may run out of space
      HINT: to retain full backups indefinitely (without warning), set option
            'repo1-retention-full' to the maximum.
```

The warning displayed is expected, since backup retention has not been configured for this example. Assuming no errors are displayed, a full backup has now been successfully created.

Restore

pgBackRest provides numerous methods and strategies for restoring a PostgreSQL database. The following section will demonstrate three forms of database restores that can be accomplished when using pgBackRest with the Crunchy Container Suite: - **Full**: restore all database files into an empty PGDATA directory - **point-in-time Recovery (PITR)**: restore a database to a specific point-in-time using an empty PGDATA directory - **Delta**: restore a database to a specific point-in-time using an existing PGDATA directory

Full This example will demonstrate a full database restore to an empty PGDATA directory. *Please ensure the Backup example is currently running and a full backup has been generated prior to running this example.*

Prior to running the full restore, we will first make a change to the currently running database, which will we will then verify still exists following the restore. Create a simple table in the database as follows:

```
$ docker exec backrest psql -c "create table backrest_test_table (id int)"
CREATE TABLE
```

Now verify that the new table exists:

```
$ docker exec backrest psql -c "table backrest_test_table"
 id
----
(0 rows)
```

With the table in place, we can now start the full restore as follows:

```
cd $CCPROOT/examples/docker/backrest/full
./run.sh
```

This will create the following in your Docker environment: - A container named **backrest-full-restore** which will perform the restore using the crunchy-backrest-restore container - A volume for the new PGDATA directory that will contain the restored database. The directory will initially be empty, as required pgBackRest when performing a full restore, and will then contain the restored database upon completion of the restore.

Please note that a brand new PV and PVC are created when running the restore to clearly indicate that the database will be restored into an entirely new (i.e. empty) volume as required by pgBackRest. The names of the new PV and PVC are as follows: - **PV**: `/${CCP_NAMESPACE}-br-new-pgdata` - **PVC**: `br-new-pgdata`

You can verify that the restore has completed successfully by verifying that the container has finished running and has exited without errors:

```
docker ps -a
```

Once the container has finished running, the post restore script can then be run, which will create a new container named **backrest-full-restored** containing the restored database:

```
cd $CCPROOT/examples/docker/backrest/full
./post-restore.sh
```

Finally, once the **backrest-full-restored** container is running we can verify that the restore was successful by verifying that the table created prior to the restore still exists:

```
$ docker exec backrest-full-restored psql -c "table backrest_test_table"
 id
----
(0 rows)
```

Please note that the default behavior of pgBackRest is to recover to the end of the WAL archive stream, which is why the full restore contained all changes made since the initial full backup was taken, including the creation of table **backrest_test_table**. pgBackRest therefore played the entire WAL archive stream for all changes that occurred up until the restore.

*As a reminder, please remember to run the cleanup script for the **Backup** example after running the cleanup script for this example.*

PITR As demonstrated with the full restore above, the default behavior of pgBackRest is to recover to the end of the WAL archive stream. However, pgBackRest also provides the ability to recover to a specific point-in-time utilizing the WAL archives created since the last backup. This example will demonstrate how pgBackRest can be utilized to perform a point-in-time recovery (PITR) and therefore recover the database to specific point-in-time specified by the user. *Please ensure that the Backup example is currently running and a full backup has been generated prior to running this example.*

Prior to running the PITR restore, we will first verify the current state of the database, after which we will then make a change to the database. This will allow us to verify that the PITR is successful by providing a method of verifying that the database has been restored to its current state following the restore.

To verify the current state of the database, we will first verify that a table called **backrest_test_table** does not exist in the database.

```
$ docker exec backrest psql -c "table backrest_test_table"
ERROR:  relation "backrest_test_table" does not exist
LINE 1:  table backrest_test_table
          ^
command terminated with exit code 1
```

Next, capture the current timestamp, which will be used later in the example when performing the restore:

```
$ docker exec backrest psql -c "select current_timestamp"
 current_timestamp
-----
2018-12-27 16:53:05.590156+00
(1 row)
```

Now create table `backrest_test_table`:

```
$ docker exec backrest psql -c "create table backrest_test_table (id int)"
CREATE TABLE
```

Then verify that the new table exists:

```
$ docker exec backrest psql -c "table backrest_test_table"
 id
----
(0 rows)
```

With the table in place, we can now start the PITR. However, the timestamp captured above must also be provided in order to instruct pgBackRest to recover to that specific point-in-time. This is done using the `CCP_BACKREST_TIMESTAMP` variable, which allows us to then start the PITR as follows (replace the timestamp in the command below with the timestamp you captured above):

```
cd $CCPROOT/examples/docker/backrest/pitr
CCP_BACKREST_TIMESTAMP="2018-12-20 09:49:02.275701+00" ./run.sh
```

This will create the following in your Docker environment: - A container named `backrest-pitr-restore` which will perform the restore using the `crunchy-backrest-restore` container

Additionally, when this example is run, the following pgBackRest environment variables are provided to the `crunchy-backrest-restore` container in order to initiate PITR to the point-in-time specified by the timestamp (in addition to any other pgBackRest variables required by the Crunchy Container Suite and pgBackRest):

```
PGBACKREST_TYPE=time
PITR_TARGET="${CCP_BACKREST_TIMESTAMP}"
```

As can be seen above, the timestamp provided for `CCP_BACKREST_TIMESTAMP` is used to populate variable `PITR_TARGET`, and therefore specify the point-in-time to restore the database to, while `PGBACKREST_TYPE` is set to `time` to indicate that a PITR should be performed.

Please note that the following pgBackRest environment variable is also set when performing the PITR, which results in a restore to a new/empty directory within an existing PV:

```
PGBACKREST_PG1_PATH=/pgdata/backrest-pitr-restored
```

You can verify that the restore has completed successfully by verifying that the container has finished running and has exited without errors:

```
docker ps -a
```

Once the container has finished running, the post restore script can then be run, which will create a new container named `backrest-pitr-restored` containing the restored database:

```
cd $CCPROOT/examples/docker/backrest/pitr
./post-restore.sh
```

Finally, once the `backrest-pitr-restored` container is running we can verify that the restore was successful by verifying that the table created prior to the restore no longer exists:

```
$ docker exec backrest-pitr-restored psql -c "table backrest_test_table"
ERROR:  relation "backrest_test_table" does not exist
LINE 1:  table backrest_test_table
          ^
command terminated with exit code 1
```

*As a reminder, please remember to run the cleanup script for the **Backup** example after running the cleanup script for this example.*

Delta By default, pgBackRest requires a clean/empty directory in order to perform a restore. However, pgBackRest also provides an another option when performing the restore in the form of the `delta` option, which allows the restore to be run against an existing PGDATA directory. With the delta option enabled, pgBackRest will use checksums to determine which files in the directory can be preserved, and which need to be restored (please note that pgBackRest will also remove any files that are not present in the backup). This example will again demonstrate a point-in-time recovery (PITR), only this time the restore will occur within the existing PGDATA directory by specifying the `delta` option during the restore. *Please ensure that the Backup example is currently running and a full backup has been generated prior to running this example.*

Prior to running the delta restore, we will first verify the current state of the database, and we will then make a change to the database. This will allow us to verify that the delta restore is successful by providing a method of verifying that the database has been restored to its current state following the restore.

To verify the current state of the database, we will first verify that a table called `backrest_test_table` does not exist in the database.

```
$ docker exec backrest psql -c "table backrest_test_table"
ERROR:  relation "backrest_test_table" does not exist
LINE 1:  table backrest_test_table
          ^
command terminated with exit code 1
```

Next, capture the current timestamp, which will be used later in the example when performing the restore:

```
$ docker exec backrest psql -c "select current_timestamp"
      current_timestamp
-----
2018-12-27 16:53:05.590156+00
(1 row)
```

Now create table **backrest_test_table**:

```
$ docker exec backrest psql -c "create table backrest_test_table (id int)"
CREATE TABLE
```

Then verify that the new table exists:

```
$ docker exec backrest psql -c "table backrest_test_table"
 id
----
(0 rows)
```

With the table in place, we can now start the delta restore. When running the restore example the timestamp captured above must also be provided in order to instruct pgBackRest to recover to that specific point-in-time. This is done using the `CCP_BACKREST_TIMESTAMP` variable, which allows us to then start the delta restore as follows (replace the timestamp in the command below with the timestamp you captured above):

```
cd $CCPROOT/examples/docker/backrest/delta
CCP_BACKREST_TIMESTAMP="2018-12-20 09:49:02.275701+00" ./run.sh
```

This will create the following in your Docker environment: - A container named **backrest-delta-restore** which will perform the restore using the `crunchy-backrest-restore` container

Additionally, when this example is run, the following pgBackRest environment variables are provided to the `crunchy-backrest-restore` container in order to initiate a delta restore to the point-in-time specified by the timestamp (in addition to any other pgBackRest variables required by the Crunchy Container Suite and pgBackRest):

```
PGBACKREST_TYPE=time
PITR_TARGET="${CCP_BACKREST_TIMESTAMP}"
PGBACKREST_DELTA=y
```

As can be seen above, the timestamp provided for `CCP_BACKREST_TIMESTAMP` is used to populate variable `PITR_TARGET`, and therefore specify the point-in-time to restore to, while `PGBACKREST_TYPE` is set to `time` to indicate that a PITR should be performed. `PGBACKREST_DELTA` is set to `y` to indicate that the delta option should be utilized when performing the restore.

It's also worth noting that the following pgBackRest environment variable is also set when performing the delta restore, which results in a restore within the existing PGDATA directory utilized by the database deployed when running the **Backup** example:

```
PGBACKREST_PG1_PATH=/pgdata/backrest
```

You can verify that the restore has completed successfully by verifying that the container has finished running and has exited without errors:

```
docker ps -a
```

Once the container has finished running, the post restore script can then be run, which will create a new container named **backrest-delta-restored** containing the restored database:

```
cd $CCPROOT/examples/docker/backrest/delta
./post-restore.sh
```

Finally, once the **backrest-delta-restored** container is running we can verify that the restore was successful by verifying that the table created prior to the restore no longer exists:

```
$ docker exec backrest-delta-restored psql -c "table backrest_test_table"
ERROR:  relation "backrest_test_table" does not exist
LINE 1:  table backrest_test_table
          ^
command terminated with exit code 1
```

As a reminder, please remember to run the *cleanup* script for the **Backup** example after running the *cleanup* script for this example.

Async Archiving

pgBackRest supports the capability to asynchronously push and get write ahead logs (WAL) to and from a WAL archive. Asynchronous archiving can improve performance by parallelizing operations, while also reducing the number of connections to remote storage. For more information on async archiving and its benefits, please see the [official pgBackRest documentation](#). This example will demonstrate how asynchronous archiving can be enabled within a crunchy-postgres or crunchy-postgres-gis container, while then also demonstrating the creation of a differential backup.

Start the example as follows:

```
cd $CCPROOT/examples/docker/backrest/async-archive
./run.sh
```

This will create the following in your Docker environment: - A container named **backrest-async-archive** containing a PostgreSQL database with pgBackRest configured - A volume for the PGDATA directory - A volume for the pgBackRest backups and archives directories

Additionally, the following variable will be set during deployment of the container in order to enable asynchronous archiving:

```
PGBACKREST_ARCHIVE_ASYNC=y
```

This will also result in the creation of the required pool path, which we can see by listing the contents of the `/pgdata` directory in the `backrest-async-archive` container:

```
$ docker exec backrest-async-archive ls /pgdata
backrest-async-archive
backrest-async-archive-backups
backrest-async-archive-spool
```

Once the database is up and running, a full backup can be taken:

```
docker exec backrest-async-archive pgbackrest backup \
  --stanza=db \
  --pg1-path=/pgdata/backrest-async-archive \
  --repo1-path=/backrestrepo/backrest-async-archive-backups \
  --log-path=/tmp \
  --type=full
```

And once a full backup has been taken, other types of backups can also be taken using pgBackRest, such as a differential backup:

```
docker exec backrest-async-archive pgbackrest backup \
  --stanza=db \
  --pg1-path=/pgdata/backrest-async-archive \
  --repo1-path=/backrestrepo/backrest-async-archive-backups \
  --log-path=/tmp \
  --type=diff
```

The following command can then be run to verify that both backups were created successfully:

```
docker exec backrest-async-archive pgbackrest info \
  --stanza=db \
  --repo1-path=/backrestrepo/backrest-async-archive-backups
```

pgBaseBackup Examples

The script assumes you are going to backup the *primary* container created in the first example, so you need to ensure that container is running. This example assumes you have configured storage as described in the link: [/installation/storage-configuration/\[Storage Configuration documentation\]](#). Things to point out with this example include its use of persistent volumes and volume claims to store the backup data files.

A successful backup will perform `pg_basebackup` on the *primary* container and store the backup in the `$CCP_STORAGE_PATH` volume under a directory named `$CCP_NAMESPACE-primary-backups`. Each backup will be stored in a subdirectory with a timestamp as the name, allowing any number of backups to be kept.

The backup script will do the following:

- Start up a backup container named `backup`
- Run `pg_basebackup` on the container named *primary*
- Store the backup in the `/tmp/backups/primary-backups` directory
- Exit after the backup

When you are ready to restore from the backup, the restore example runs a PostgreSQL container using the backup location. Upon initialization, the container will use `rsync` to copy the backup data to this new container and then launch PostgreSQL using the original backed-up data.

The restore script will do the following:

- Start up a container named *restore*
- Copy the backup files from the previous backup example into `/pgdata`
- Start up the container using the backup files
- Map the PostgreSQL port of 5432 in the container to your local host port of 12001

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

Run the backup with this command:

```
cd $CCPROOT/examples/docker/backup
./run.sh
```

When you're ready to restore, a *restore* example is provided.

It's required to specify a backup path for this example. To get the correct path check the `backup` job logs or a timestamp:

```
docker logs backup-vpk91 | grep BACKUP_PATH
Wed May 9 20:32:00 UTC 2018 INFO: BACKUP_PATH is set to
/pgdata/primary-backups/2018-05-09-20-32-00.
```

`BACKUP_PATH` can also be discovered by looking at the backup mount directly (if access to the storage is available to the user).

An example of `BACKUP_PATH` is as followed:

```
"name": "BACKUP_PATH",
"value": "primary-backups/2018-05-09-20-32-00"
```

When you are ready to restore from the backup created, run the following example:

```
cd $CCPROOT/examples/docker/restore
./run.sh
```

Kubernetes and OpenShift

Running the example:

```
cd $CCPROOT/examples/kube/backup
./run.sh
```

The Kubernetes Job type executes a pod and then the pod exits. You can view the Job status using this command:

```
${CCP_CLI} get job
```

When you're ready to restore, a *restore* example is provided.

It's required to specify a backup path for this example. To get the correct path check the `backup` job logs or a timestamp:

```
kubect1 logs backup-vpk91 | grep BACKUP_PATH
Wed May 9 20:32:00 UTC 2018 INFO: BACKUP_PATH is set to
/pgdata/primary-backups/2018-05-09-20-32-00.
```

`BACKUP_PATH` can also be discovered by looking at the backup mount directly (if access to the storage is available to the user).

An example of `BACKUP_PATH` defined as a variable within the JSON script is as follows:

```
"name": "BACKUP_PATH",
"value": "primary-backups/2018-05-09-20-32-00"
```

Running the example:

```
cd $CCPROOT/examples/kube/restore
./run.sh
```

Test the restored database as follows:

```
psql -h restore -U postgres postgres
```

Point in Time Recovery (PITR)

PITR (point-in-time-recovery) is a feature that allows for recreating a database from backup and log files at a certain point in time. This is done using a write ahead log (WAL) which is kept in the `pg_wal` directory within `PGDATA`. Changes made to the database files over time are recorded in these log files, which allows it to be used for disaster recovery purposes.

When using PITR as a backup method, in order to restore from the last checkpoint in the event of a database or system failure, it is only necessary to save these log files plus a full backup. This provides an additional advantage in that it is not necessary to keep multiple full backups on hand, which consume space and time to create. This is because point in time recovery allows you to “replay” the log files and recover your database to any point since the last full backup.

More detailed information about Write Ahead Log (WAL) archiving can be found [link:https://www.postgresql.org/docs/10/static/continuous-archiving.html](https://www.postgresql.org/docs/10/static/continuous-archiving.html)[here.]

By default in the `crunchy-postgres` container, WAL logging is *not* enabled. To enable WAL logging *outside of this example*, set the following environment variables when starting the `crunchy-postgres` container:

```
ARCHIVE_MODE=on
ARCHIVE_TIMEOUT=60
```

These variables set the same name settings within the `postgresql.conf` file that is used by the database. When set, WAL files generated by the database will be written out to the `/pgwal` mount point.

A full backup is required to do a PITR. `crunchy-backup` currently performs this role within the example, running a `pg_basebackup` on the database. This is a requirement for PITR. After a backup is performed, code is added into `crunchy-postgres` which will also check to see if you want to do a PITR.

There are three volume mounts used with the PITR example.

- `/recover` - When specified within a `crunchy-postgres` container, PITR is activated during container startup.
- `/backup` - This is used to find the base backup you want to recover from.
- `/pgwal` - This volume is used to write out new WAL files from the newly restored database container.

Some environment variables used to manipulate the point in time recovery logic:

- The `RECOVERY_TARGET_NAME` environment variable is used to tell the PITR logic what the name of the target is.
- `RECOVERY_TARGET_TIME` is also an optional environment variable that restores using a known time stamp.

If you don't specify either of these environment variables, then the PITR logic will assume you want to restore using all the WAL files or essentially the last known recovery point.

The `RECOVERY_TARGET_INCLUSIVE` environment variable is also available to let you control the setting of the `recovery.conf` setting `recovery_target_inclusive`. If you do not set this environment variable the default is `true`.

Once you recover a database using PITR, it will be in read-only mode. To make the database resume as a writable database, run the following SQL command:

```
postgres=# select pg_wal_replay_resume();
```

If you're running the PITR example for *PostgreSQL versions 9.5 or 9.6*, please note that starting in PostgreSQL version 10, the `pg_xlog` directory was renamed to `pg_wal`. Additionally, all usages of the function `pg_xlog_replay_resume` were changed to `pg_wal_replay_resume`.

It takes about 1 minute for the database to become ready for use after initially starting.

WAL segment files are written to the `/tmp` directory. Leaving the example running for a long time could fill up your `/tmp` directory.

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

Create a database container as follows:

```
cd $CCPROOT/examples/docker/pitr
./run-pitr.sh
```

Next, we will create a base backup of that database using this:

```
./run-backup-pitr.sh
```

After creating the base backup of the database, WAL segment files are created every 60 seconds that contain any database changes. These segments are stored in the `/tmp/pitr/pitr/pg_wal` directory.

Next, create some recovery targets within the database by running the SQL commands against the `pitr` database as follows:

```
./run-sql.sh
```

This will create recovery targets named `beforechanges`, `afterchanges`, and `nomorechanges`. It will create a table, `pitrtest`, between the `beforechanges` and `afterchanges` targets. It will also run a SQL `CHECKPOINT` to flush out the changes to WAL segments. These labels can be used to mark the points in the recovery process that will be referenced when creating the restored database.

Next, now that we have a base backup and a set of WAL files containing our database changes, we can shut down the `pitr` database to simulate a database failure. Do this by running the following:

```
docker stop pitr
```

Next, let's edit the restore script to use the base backup files created in the step above. You can view the backup path name under the `/tmp/backups/pitr-backups/` directory. You will see another directory inside of this path with a name similar to `2018-03-21-21-03-29`. Copy and paste that value into the `run-restore-pitr.sh` script in the `BACKUP` environment variable.

After that, run the script.

```
vi ./run-restore-pitr.sh
./run-restore-pitr.sh
```

The WAL segments are read and applied when restoring from the database backup. At this point, you should be able to verify that the database was restored to the point before creating the test table:

```
psql -h 127.0.0.1 -p 12001 -U postgres postgres -c 'table pitrtest'
```

This SQL command should show that the `pitrtest` table does not exist at this recovery time. The output should be similar to:

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

PostgreSQL allows you to pause the recovery process if the target name or time is specified. This pause would allow a DBA a chance to review the recovery time/name and see if this is what they want or expect. If so, the DBA can run the following command to resume and complete the recovery:

```
psql -h 127.0.0.1 -p 12001 -U postgres postgres -c 'select pg_wal_replay_resume()'
```

Until you run the statement above, the database will be left in read-only mode.

Next, run the script to restore the database to the `afterchanges` restore point. Update the `RECOVERY_TARGET_NAME` to `afterchanges`:

```
vi ./run-restore-pitr.sh
./run-restore-pitr.sh
```

After this restore, you should be able to see the test table:

```
psql -h 127.0.0.1 -p 12001 -U postgres postgres -c 'table pitrtest'
psql -h 127.0.0.1 -p 12001 -U postgres postgres -c 'select pg_wal_replay_resume()'
```

Lastly, start a recovery using all of the WAL files. This will get the restored database as current as possible. To do so, edit the script to remove the `RECOVERY_TARGET_NAME` environment setting completely:

```
./run-restore-pitr.sh
sleep 30
psql -h 127.0.0.1 -p 12001 -U postgres postgres -c 'table pitrtest'
psql -h 127.0.0.1 -p 12001 -U postgres postgres -c 'create table foo (id int)'
```

At this point, you should be able to create new data in the restored database and the test table should be present. When you recover the entire WAL history, resuming the recovery is not necessary to enable writes.

Kubernetes and OpenShift

Start by running the example database container:

```
cd $CCPROOT/examples/kube/pitr
./run-pitr.sh
```

This step will create a database container, *pitr*. This container is configured to continuously write WAL segment files to a mounted volume (*/pgwal*).

After you start the database, you will create a base backup using this command:

```
./run-backup-pitr.sh
```

This will create a backup and write the backup files to a persistent volume (*/pgbackup*).

Next, create some recovery targets within the database by running the SQL commands against the *pitr* database as follows:

```
./run-sql.sh
```

This will create recovery targets named **beforechanges**, **afterchanges**, and **nomorechanges**. It will create a table, *pitrtest*, between the **beforechanges** and **afterchanges** targets. It will also run a SQL CHECKPOINT to flush out the changes to WAL segments.

Next, now that we have a base backup and a set of WAL files containing our database changes, we can shut down the *pitr* database to simulate a database failure. Do this by running the following:

```
/${CCP_CLI} delete pod pitr
```

Next, we will create 3 different restored database containers based upon the base backup and the saved WAL files.

First, get the `BACKUP_PATH` created by the `backup-pitr` example by viewing the pods logs:

```
/${CCP_CLI} logs backup-pitr-8sfkh | grep PATH
Thu May 10 18:07:58 UTC 2018 INFO: BACKUP_PATH is set to /pgdata/pitr-backups/2018-05-10-18-07-58.
```

Edit the `restore-pitr.json` file and change the `BACKUP_PATH` environment variable using the path discovered above (note: `/pgdata/` is not required and should be excluded in the variable):

```
{
  "name": "BACKUP_PATH",
  "value": "pitr-backups/2018-05-10-18-07-58"
}
```

Next, we restore prior to the **beforechanges** recovery target. This recovery point is *before* the *pitrtest* table is created.

Edit the `restore-pitr.json` file, and edit the environment variable to indicate we want to use the **beforechanges** recovery point:

```
{
  "name": "RECOVERY_TARGET_NAME",
  "value": "beforechanges"
}
```

Then run the following to create the restored database container:

```
./run-restore-pitr.sh
```

After the database has restored, you should be able to perform a test to see if the recovery worked as expected:

```
psql -h restore-pitr -U postgres postgres -c 'table pitrtest'
psql -h restore-pitr -U postgres postgres -c 'create table foo (id int)'
```

The output of these commands should show that the *pitrtest* table is not present. It should also show that you can not create a new table because the database is paused in read-only mode.

To make the database resume as a writable database, run the following SQL command:

```
select pg_wal_replay_resume();
```

It should then be possible to write to the database:

```
psql -h restore-pitr -U postgres postgres -c 'create table foo (id int)'
```

You can also test that if **afterchanges** is specified, that the *pitrtest* table is present but that the database is still in recovery mode.

Lastly, you can test a full recovery using *all* of the WAL files, if you remove the `RECOVERY_TARGET_NAME` environment variable completely.

The storage portions of this example can all be found under `$CCP_STORAGE_PATH/$CCP_NAMESPACE-restore-p`

title: “pgDump” date: draft: false weight: 32 —

pg_dump example

The script assumes you are going to backup the *primary* example and that container is running.

This example assumes you have configured a storage filesystem as described in the link:[/installation/storage-configuration/Storage Configuration](#) document.

A successful backup will perform `pg_dump/pg_dumpall` on the primary and store the resulting files in the mounted volume under a directory named `<HOSTNAME>-backups` as a sub-directory, then followed by a unique backup directory based upon a date and timestamp - allowing any number of backups to be kept.

For more information on how to configure this container, please see the link:[/container-specifications/\[Container Specifications\]](#) document.

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

Run the backup with this command:

```
cd $CCPROOT/examples/docker/pgdump
./run.sh
```

Kubernetes and OpenShift

Running the example:

```
cd $CCPROOT/examples/kube/pgdump
./run.sh
```

The Kubernetes Job type executes a pod and then the pod exits. You can view the Job status using this command:

```
`${CCP_CLI} get job
```

The `pgdump.json` file within that directory specifies options that control the behavior of the `pgdump` job. Examples of this include whether to run `pg_dump` vs `pg_dumpall` and advanced options for specific backup use cases.

pg_restore example

The script assumes you are going to restore to the *primary* example and that container is running and a backup has been created using the `pgdump` example..

This example assumes you have configured a storage filesystem as described in the link:[/installation/storage-configuration/Storage Configuration](#) document.

Successful use of the `crunchy-pgrestore` container will run a job to restore files generated by `pg_dump/pg_dumpall` to a container via `psql/pg_restore`; then container will terminate successfully and signal job completion.

For more information on how to configure this container, please see the link:[/container-specifications/\[Container Specifications\]](#) document.

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

Run the restore with this command:

```
cd $CCPROOT/examples/docker/pgrestore
./run.sh
```

Kubernetes and OpenShift

By default, pgrestore container will automatically restore from the most recent backup. If you want to restore to a specific backup, edit the `pgrestore.json` file and update the `PGRESTORE_BACKUP_TIMESTAMP` setting to specify the backup path you want to restore with. For example:

```
"name": "PGRESTORE_BACKUP_TIMESTAMP",  
"value": "2018-03-27-14-35-33"
```

Running the example:

```
cd $CCPROOT/examples/kube/pgrestore  
./run.sh
```

The Kubernetes Job type executes a pod and then the pod exits. You can view the Job status using this command:

```
`${CCP_CLI} get job
```

The `pgrestore.json` file within that directory specifies options that control the behavior of the pgrestore job.

Crunchy Scheduler

The Crunchy Scheduler container implements a cronlike microservice within a namespace to automate backups of a PostgreSQL database.

Currently Crunchy Scheduler only supports two types of tasks:

- `pgBackRest`
- `pgBaseBackup`

This service watches Kubernetes for config maps with the label `crunchy-scheduler=true`. If found the scheduler will parse the data found in the config map (json object) and convert it to a scheduled task. If the config map is removed, the scheduler will delete the task.

See the following examples for creating config maps that Crunchy Scheduler can parse:

- link:<https://github.com/CrunchyData/crunchy-containers/blob/master/examples/kube/scheduler/configs/schedule-backrest-diff.json>[pgBackRest Diff Backup]
- link:<https://github.com/CrunchyData/crunchy-containers/blob/master/examples/kube/scheduler/configs/schedule-backrest-full.json>[pgBackRest Full Backup]
- link:<https://github.com/CrunchyData/crunchy-containers/blob/master/examples/kube/scheduler/configs/schedule-pgbasebackup.json>[Backup]

The Crunchy Scheduler requires a Service Account to create jobs (`pgBaseBackup`) and to exec (`pgBackRest`). See the link:<https://github.com/CrunchyData/crunchy-containers/blob/master/examples/kube/scheduler/scheduler-sa.json>[scheduler example] for the required permissions on this account.

Crunchy Scheduler uses the UTC timezone for all schedules.

pgBackRest Schedules

To configure Crunchy Scheduler to create `pgBackRest` backups the following is required:

- `pgBackRest` schedule definition requires a deployment name. The PostgreSQL pod should be created by a deployment.

pgBaseBackup Schedules

To configure Crunchy Scheduler to create `pgBaseBackup` scheduled backups, the following is required:

- The name of the secret that contains the username and password the Scheduler will use to configure the job template. See link:<https://github.com/CrunchyData/crunchy-containers/blob/master/examples/kube/scheduler/primary/secret.json>[the primary secret example]. for the structure required by the Scheduler.
- The name of the PVC created for the backups. This should be created by the user prior to scheduling the task.

When using `pgBaseBackup` schedules, it may be required to apply specific `supplementalGroups` or an `fsGroup` to the backup job created by the scheduler. To apply a specific `securityContext` for your storage provider, mount a `backup-template.json` to `/configs` on the scheduler pod.

For an example of applying a custom template, link:<https://github.com/CrunchyData/crunchy-containers/blob/master/examples/kube/scheduler/primary/secret.json>[the schedule example].

Kubernetes and OpenShift

First, start the PostgreSQL example created for the Scheduler by running the following commands:

```
# Kubernetes
cd $CCPROOT/examples/kube/scheduler/primary
./run.sh
```

The pod created should show a ready status before proceeding.

Next, start the scheduler by running the following command:

```
# Kubernetes
cd $CCPROOT/examples/kube/scheduler
./run.sh
```

Once the scheduler is deployed, register the backup tasks by running the following command:

```
# Kubernetes
cd $CCPROOT/examples/kube/scheduler
./add-schedules.sh
```

The scheduled tasks will (these are just for fast results, not recommended for production):

- take a backup every minute using pgBaseBackup
- take a full pgBackRest backup every even minute
- take a diff pgBackRest backup every odd minute

View the logs for the scheduler pod until the tasks run:

```
`${CCP_CLI?}` logs scheduler -f
```

View the pgBaseBackup pods results after the backup completes:

```
`${CCP_CLI?}` logs <basebackup pod name>
```

View the pgBackRest backups via exec after the backup completes:

```
`${CCP_CLI?}` exec -ti <primary deployment pod name> -- pgbackrest info
```

Clean up the examples by running the following commands:

```
$CCPROOT/examples/kube/scheduler/primary/cleanup.sh
$CCPROOT/examples/kube/scheduler/cleanup.sh
```

Centralized Logging Example

The logs generated by containers are critical for deployments because they provide insights into the health of the system. PostgreSQL logs are very detailed and there is some information that can only be obtained from logs (but not limited to):

- Connections and Disconnections of users
- Checkpoint Statistics
- PostgreSQL Server Errors

Aggregating container logs across multiple hosts allows administrators to audit, debug problems and prevent repudiation of misconduct.

In the following example we will demonstrate how to setup Kubernetes and OpenShift to use centralized logging by using an EFK (Elasticsearch, Fluentd and Kibana) stack. Fluentd will run as a daemonset on each host within the Kubernetes cluster and extract container logs, Elasticsearch will consume and index the logs gathered by Fluentd and Kibana will allow users to explore and visualize the logs via a web dashboard.

To learn more about the EFK stack, see the following:

- <https://www.elastic.co/products/elasticsearch>
- <https://www.fluentd.org/architecture>
- <https://www.elastic.co/products/kibana>

Configure PostgreSQL for Centralized Logging

By default, Crunchy PostgreSQL logs to files in the `/pgdata` directory. In order to get the logs out of the container we need to configure PostgreSQL to log to `stdout`.

The following settings should be configured in `postgresql.conf` to make PostgreSQL log to `stdout`:

```
log_destination = 'stderr'  
logging_collector = off
```

Changes to logging settings require a restart of the PostgreSQL container to take effect.

Deploying the EFK Stack On OpenShift Container Platform

OpenShift Container Platform can be installed with an EFK stack. For more information about configuring OpenShift to create an EFK stack, see the official documentation:

- https://docs.openshift.com/container-platform/3.11/install_config/aggregate_logging.html

Deploying the EFK Stack On Kubernetes

First, deploy the EFK stack by running the example using the following commands:

```
cd $CCPROOT/examples/kube/centralized-logging/efk  
./run.sh
```

Elasticsearch is configured to use an `emptyDir` volume in this example. Configure this example to provide a persistent volume when deploying into production.

Next, verify the pods are running in the `kube-system` namespace:

```
/${CCP_CLI?} get pods -n kube-system --selector=k8s-app=elasticsearch-logging  
/${CCP_CLI?} get pods -n kube-system --selector=k8s-app=fluentd-es  
/${CCP_CLI?} get pods -n kube-system --selector=k8s-app=kibana-logging
```

If all pods deployed successfully, Elasticsearch should already be receiving container logs from Fluentd.

Next we will deploy a PostgreSQL Cluster (primary and replica deployments) to demonstrate PostgreSQL logs are being captured by Fluentd.

Deploy the PostgreSQL cluster by running the following:

```
cd $CCPROOT/examples/kube/centralized-logging/postgres-cluster  
./run.sh
```

Next, verify the pods are running:

```
/${CCP_CLI?} get pods --selector=k8s-app=postgres-cluster
```

With the PostgreSQL successfully deployed, we can now query the logs in Kibana.

We will need to setup a port-forward to the Kibana pod to access it. To do that we first get the name of the pod by running the following command:

```
/${CCP_CLI?} get pod --selector=k8s-app=kibana-logging -n kube-system
```

Next, start the port-forward:

```
/${CCP_CLI?} port-forward <KIBANA POD NAME> 5601:5601 -n kube-system
```

To access the web dashboard navigate in a browser to `127.0.0.1:5601`.

First, click the **Discover** tab and setup an index pattern to use for queries.

The index pattern name we will use is `logstash-*` because Fluentd is configured to generate logstash style logs.

Next we will configure the **Time Filter** field name to be `@timestamp`.

Now that our index pattern is created, we can query for the container logs.

Click the **Discover** tab and use the following queries:

```
# KUBERNETES
CONTAINER_NAME: *primary* AND MESSAGE: ".*LOG*"
# OpenShift
kubernetes.pod_name: "primary" AND log
```

For more information about querying Kibana, see the official documentation: <https://www.elastic.co/guide/en/beats/packetbeat/current/kibana-queries-filters.html>

pgAudit Enhanced Logging

This example provides an example of enabling `pg_audit` output. As of release 1.3, `pg_audit` is included in the `crunchy-postgres` container and is added to the PostgreSQL shared library list in `postgresql.conf`.

Given the numerous ways `pg_audit` can be configured, the exact `pg_audit` configuration is left to the user to define. `pg_audit` allows you to configure auditing rules either in `postgresql.conf` or within your SQL script.

For this test, we place `pg_audit` statements within a SQL script and verify that auditing is enabled and working. If you choose to configure `pg_audit` via a `postgresql.conf` file, then you will need to define your own custom file and mount it to override the default `postgresql.conf` file.

Docker

Run the following to create a database container:

```
cd $CCPROOT/examples/docker/pgaudit
./run.sh
```

This starts an instance of the `pg_audit` container (running `crunchy-postgres`) on port 12005 on localhost. The test script is then automatically executed.

This test executes a SQL file which contains `pg_audit` configuration statements as well as executes some basic SQL commands. These SQL commands will cause `pg_audit` to create log messages in the `pg_log` log file created by the database container.

Kubernetes and OpenShift

Run the following:

```
cd $CCPROOT/examples/kube/pgaudit
./run.sh
```

This script will create a PostgreSQL pod with the `pgAudit` extension configured and ready to use

Once the pod is deployed successfully run the following command to test the extension:

```
cd $CCPROOT/examples/kube/pgaudit
./test-pgaudit.sh
```

This example has been configured to log directly to stdout of the pod. To view the PostgreSQL logs, run the following:

```
$CCP_CLI logs pgbaudit
```

Metrics Collection

You can collect various PostgreSQL metrics from your database container by running a `crunchy-collect` container that points to your database container.

This example starts up 5 containers:

- Collect (`crunchy-collect`)
- Grafana (`crunchy-grafana`)
- PostgreSQL (`crunchy-postgres`)
- Prometheus (`crunchy-prometheus`)

Every 5 seconds by default, Prometheus will scrape the Collect container for metrics. These metrics will then be visualized by Grafana, which by default can be accessed with the following credentials:

- Username : *admin*
- Password: *password*

By default, Prometheus detects which environment its running on (Docker, Kubernetes, or OpenShift Container Platform) and applies a default configuration.

When running in Kuberenetes and OpenShift, the following two labels are required by the deployments:

- "crunchy_collect": "true"
- "name": "some-pod-name-here"

The `crunchy_collect` label allows Prometheus to find all pods that are serving metrics to be scraped for storage.

The `name` label allows Prometheus to rewrite the name of the pod so if it changes there's not duplicate entries.

Additionally, the collect container uses a special PostgreSQL role `ccp_monitoring`. This user is created by setting the `PGMONITOR_PASSWORD` environment variable on the PostgreSQL container.

Discovering pods requires a cluster role service account. See the link:<https://github.com/crunchydata/crunchy-containers/blob/master/examples/and-OpenShift/metrics-JSON-file-for-more-details>.

For Docker environments the collect hostname must be specified as an environment variable.

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

To start this set of containers, run the following:

```
cd $CCPROOT/examples/docker/metrics
./run.sh
```

You will be able to access the Grafana and Prometheus services from the following web addresses:

- Grafana (<http://0.0.0.0:3000>)
- Prometheus (<http://0.0.0.0:9090>)

The crunchy-postgres container is accessible on port `5432`.

Kubernetes and OpenShift

Running the example:

```
cd $CCPROOT/examples/kube/metrics
./run.sh
```

It's required to use `port-forward` to access the Grafana dashboard. To start the port-forward, run the following command:

```
${CCP_CLI} port-forward metrics 3000:3000
${CCP_CLI} port-forward metrics 9090:9090
```

- Grafana dashboard can be then accessed from <http://127.0.0.01:3000>
- Prometheus dashboard can be then accessed from <http://127.0.0.01:9090>

You can view the container logs using these command:

```
${CCP_CLI} logs -c grafana metrics
${CCP_CLI} logs -c prometheus metrics
${CCP_CLI} logs -c collect primary-metrics
${CCP_CLI} logs -c postgres primary-metrics
${CCP_CLI} logs -c collect replica-metrics
${CCP_CLI} logs -c postgres replica-metrics
```

pgBadger Example

pgbadger is a PostgreSQL tool that reads the log files from a specified database in order to produce a HTML report that shows various PostgreSQL statistics and graphs. This example runs the pgbadger HTTP server against a crunchy-postgres container and illustrates how to view the generated reports.

The port utilized for this tool is port 14000 for Docker environments and port 10000 for Kubernetes and OpenShift environments.

The container creates a default database called *userdb*, a default user called *testuser* and a default password of *password*.

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

Run the example as follows:

```
cd $CCPROOT/examples/docker/pgbadger
./run.sh
```

After execution, the container will run and provide a simple HTTP command you can browse to view the report. As you run queries against the database, you can invoke this URL to generate updated reports:

```
curl -L http://127.0.0.1:14000/api/badgergenerate
```

Kubernetes and OpenShift

Running the example:

```
cd $CCPROOT/examples/kube/pgbadger
./run.sh
```

After execution, the container will run and provide a simple HTTP command you can browse to view the report. As you run queries against the database, you can invoke this URL to generate updated reports:

```
curl -L http://pgbadger:10000/api/badgergenerate
```

You can view the database container logs using these commands:

```
${CCP_CLI} logs pgbadger -c pgbadger
${CCP_CLI} logs pgbadger -c postgres
```

pgBouncer Connection Pooling Example

Crunchy pgBouncer is a lightweight connection pooler for PostgreSQL databases.

The following examples create the following containers:

- pgBouncer Primary
- pgBouncer Replica
- PostgreSQL Primary
- PostgreSQL Replica

In Kubernetes and OpenShift, this example will also create:

- pgBouncer Primary Service
- pgBouncer Replica Service
- Primary Service
- Replica Service
- PostgreSQL Secrets
- pgBouncer Secrets

To cleanup the objects created by this example, run the following in the **pgbouncer** example directory:

```
./cleanup.sh
```

For more information on **pgBouncer**, see the link:<https://pgbouncer.github.io>[official website].

Docker

Run the pgbouncer example:

```
cd $CCPROOT/examples/docker/pgbouncer
./run.sh
```

Once all containers have deployed and are ready for use, `psql` to the target databases through `pgBouncer`:

```
psql -d userdb -h 0.0.0.0 -p 6432 -U testuser
psql -d userdb -h 0.0.0.0 -p 6433 -U testuser
```

To connect to the administration database within `pgbouncer`, connect using `psql`:

```
psql -d pgbouncer -h 0.0.0.0 -p 6432 -U pgbouncer
psql -d pgbouncer -h 0.0.0.0 -p 6433 -U pgbouncer
```

Kubernetes and OpenShift

OpenShift: If custom configurations aren't being mounted, an `emptydir` volume is required to be mounted at `/pgconf`.

Run the pgbouncer example:

```
cd $CCPROOT/examples/kube/pgbouncer
./run.sh
```

Once all containers have deployed and are ready for use, `psql` to the target databases through `pgBouncer`:

```
psql -d userdb -h pgbouncer-primary -p 6432 -U testuser
psql -d userdb -h pgbouncer-replica -p 6432 -U testuser
```

To connect to the administration database within `pgbouncer`, connect using `psql`:

```
psql -d pgbouncer -h pgbouncer-primary -p 6432 -U pgbouncer -c "SHOW SERVERS"
psql -d pgbouncer -h pgbouncer-replica -p 6432 -U pgbouncer -c "SHOW SERVERS"
```

pgPool Logical Router Example

An example is provided that will run a `pgPool II` container in conjunction with the `primary-replica` example provided above.

You can execute both `INSERT` and `SELECT` statements after connecting to `pgpool`. The container will direct `INSERT` statements to the primary and `SELECT` statements will be sent round-robin to both the primary and replica.

The container creates a default database called `userdb`, a default user called `testuser` and a default password of `password`.

You can view the nodes that `pgpool` is configured for by running:

```
psql -h pgpool -U testuser userdb -c 'show pool_nodes'
```

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

Create the container as follows:

```
cd $CCPROOT/examples/docker/pgpool
./run.sh
```

The example is configured to allow the `testuser` to connect to the `userdb` database.

```
psql -h localhost -U testuser -p 12003 userdb
```

Kubernetes and OpenShift

Run the following command to deploy the pgpool service:

```
cd $CCPROOT/examples/kube/pgpool
./run.sh
```

The example is configured to allow the *testuser* to connect to the *userdb* database.

```
psql -h pgpool -U testuser userdb
```

SSL Authentication

This example shows how you can configure PostgreSQL to use SSL for client authentication.

The example requires SSL certificates and keys to be created. Included in the examples directory is a script to create self-signed certificates (server and client) for the example: `$CCPROOT/examples/ssl-creator.sh`.

The example creates a client certificate for the user *testuser*. Furthermore, the server certificate is created for the server name *custom-config-ssl*.

This example can be run as follows for the Docker environment:

```
cd $CCPROOT/examples/docker/custom-config-ssl
./run.sh
```

And the example can be run in the following directory for the Kubernetes and OpenShift environments:

```
cd $CCPROOT/examples/kube/custom-config-ssl
./run.sh
```

A required step to make this example work is to define in your `/etc/hosts` file an entry that maps *custom-config-ssl* to the service IP address for the container.

For instance, if your service has an address as follows:

```
/${CCP_CLI} get service
NAME                CLUSTER-IP          EXTERNAL-IP          PORT(S)              AGE
custom-config-ssl  172.30.211.108      <none>               5432/TCP
```

Then your `/etc/hosts` file needs an entry like this:

```
172.30.211.108 custom-config-ssl
```

For production Kubernetes and OpenShift installations, it will likely be preferred for DNS names to resolve to the PostgreSQL service name and generate server certificates using the DNS names instead of the example name *custom-config-ssl*.

If as a client it's required to confirm the identity of the server, *verify-full* can be specified for *ssl-mode* in the connection string. This will check if the server and the server certificate have the same name. Additionally, the proper connection parameters must be specified in the connection string for the certificate information required to trust and verify the identity of the server (*sslrootcert* and *sslcr1*), and to authenticate the client using a certificate (*sslcert* and *sslkey*):

```
psql "postgresql://testuser@custom-config-ssl:5432/userdb?\
sslmode=verify-full&\
sslrootcert=$CCPROOT/examples/kube/custom-config-ssl/certs/ca.crt&\
sslcr1=$CCPROOT/examples/kube/custom-config-ssl/certs/ca.crl&\
sslcert=$CCPROOT/examples/kube/custom-config-ssl/certs/client.crt&\
sslkey=$CCPROOT/examples/kube/custom-config-ssl/certs/client.key"
```

To connect via IP, *sslmode* can be changed to *require*. This will verify the server by checking the certificate chain up to the trusted certificate authority, but will not verify that the hostname matches the certificate, as occurs with *verify-full*. The same connection parameters as above can be then provided for the client and server certificate information.

```
psql "postgresql://testuser@IP_OF_PGSQL:5432/userdb?\
sslmode=require&\
sslrootcert=$CCPROOT/examples/kube/custom-config-ssl/certs/ca.crt&\
sslcr1=$CCPROOT/examples/kube/custom-config-ssl/certs/ca.crl&\
sslcert=$CCPROOT/examples/kube/custom-config-ssl/certs/client.crt&\
sslkey=$CCPROOT/examples/kube/custom-config-ssl/certs/client.key"
```

You should see a connection that looks like the following:

```
psql (10.6)
SSL connection (protocol: TLSv1.2, cipher: ECDHE-RSA-AES256-GCM-SHA384, bits: 256, compression:
off)
Type "help" for help.

userdb=>
```

Custom Configuration

You can use your own version of the SQL file `setup.sql` to customize the initialization of database data and objects when the container and database are created.

This works by placing a file named `setup.sql` within the `/pgconf` mounted volume directory. Portions of the `setup.sql` file are required for the container to work; please see comments within the sample `setup.sql` file.

If you mount a `/pgconf` volume, `crunchy-postgres` will look at that directory for `postgresql.conf`, `pg_hba.conf`, `pg_ident.conf`, SSL server/ca certificates and `setup.sql`. If it finds one of them it will use that file instead of the default files.

Docker

This example can be run as follows for the Docker environment:

```
cd $CCPROOT/examples/docker/custom-config
./run.sh
```

Kubernetes and OpenShift

Running the example:

```
cd $CCPROOT/examples/kube/custom-config
./run.sh
```

Geospatial (PostGIS)

An example is provided that will run a PostgreSQL with PostGIS pod and service in Kubernetes and OpenShift and a container in Docker. The container creates a default database called `userdb`, a default user called `testuser` and a default password of `password`.

You can view the extensions that `postgres-gis` has enabled by running the following command and viewing the listed PostGIS packages:

```
psql -h postgres-gis -U testuser userdb -c '\dx'
```

To validate that PostGIS is installed and which version is running, run the command:

```
psql -h postgres-gis -U testuser userdb -c "SELECT postgis_full_version();"
```

You should expect to see output similar to:

```
postgis_full_version
-----
POSTGIS="2.4.2 r16113" PGSQL="100" GEOS="3.5.0-CAPI-1.9.0 r4084" PROJ="Rel. 4.8.0, 6 March 2012"
GDAL="GDAL 1.11.4, released 2016/01/25" LIBXML="2.9.1" LIBJSON="0.11" TOPOLOGY RASTER
(1 row)
```

As an exercise for invoking some of the basic PostGIS functionality for validation, try defining a 2D geometry point while giving inputs of longitude and latitude through this command.

```
psql -h postgres-gis -U testuser userdb -c "select ST_MakePoint(28.385200, -81.563900);"
```

You should expect to see output similar to:

```
st_makepoint
-----
01010000000516B9A779C623C40B98D06F0166454C0
(1 row)
```

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

Create the container as follows:

```
cd $CCPROOT/examples/docker/postgres-gis
./run.sh
```

Enter the following command to connect to the postgres-gis container that is mapped to your local port 12000:

```
psql -h localhost -U testuser -p 12000 userdb
```

Kubernetes and OpenShift

Running the example:

```
cd $CCPROOT/examples/kube/postgres-gis
./run.sh
```

Replication

This example starts a primary and a replica pod containing a PostgreSQL database.

The container creates a default database called *userdb*, a default user called *testuser* and a default password of *password*.

For the Docker environment, the script additionally creates:

- A docker volume using the local driver for the primary
- A docker volume using the local driver for the replica
- A container named *primary* binding to port 12007
- A container named *replica* binding to port 12008
- A mapping of the PostgreSQL port 5432 within the container to the localhost port 12000
- The database using predefined environment variables

And specifically for the Kubernetes and OpenShift environments:

- emptyDir volumes for persistence
- A pod named *pr-primary*
- A pod named *pr-replica*
- A pod named *pr-replica-2*
- A service named *pr-primary*
- A service named *pr-replica*
- The database using predefined environment variables

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

To create the example and run the container:

```
cd $CCPROOT/examples/docker/primary-replica
./run.sh
```

Connect from your local host as follows:

```
psql -h localhost -p 12007 -U testuser -W userdb
psql -h localhost -p 12008 -U testuser -W userdb
```

Kubernetes and OpenShift

Run the following command to deploy a primary and replica database cluster:

```
cd $CCPROOT/examples/kube/primary-replica
./run.sh
```

It takes about a minute for the replica to begin replicating with the primary. To test out replication, see if replication is underway with this command:

```
`${CCP_CLI?}` exec -ti pr-primary -- psql -d postgres -c 'table pg_stat_replication'
```

If you see a line returned from that query it means the primary is replicating to the replica. Try creating some data on the primary:

```
`${CCP_CLI?}` exec -ti pr-primary -- psql -d postgres -c 'create table foo (id int)'
`${CCP_CLI?}` exec -ti pr-primary -- psql -d postgres -c 'insert into foo values (1)'
```

Then verify that the data is replicated to the replica:

```
`${CCP_CLI?}` exec -ti pr-replica -- psql -d postgres -c 'table foo'
```

primary-replica-dc

If you wanted to experiment with scaling up the number of replicas, you can run the following example:

```
cd $CCPROOT/examples/kube/primary-replica-dc
./run.sh
```

You can verify that replication is working using the same commands as above.

```
`${CCP_CLI?}` exec -ti primary-dc -- psql -d postgres -c 'table pg_stat_replication'
```

Helm

This example resides under the `$CCPROOT/examples/helm` directory. View the README to run this example using Helm link:<https://github.com/CrunchyData/crunchy-containers/blob/master/examples/helm/primary-replica/README.md>[here].

PostgreSQL Container Example

This example starts a single PostgreSQL container and service, the most simple of examples.

The container creates a default database called *userdb*, a default user called *testuser* and a default password of *password*.

For all environments, the script additionally creates:

- A persistent volume claim
- A crunchy-postgres container named *primary*
- The database using predefined environment variables

And specifically for the Kubernetes and OpenShift environments:

- A pod named *primary*
- A service named *primary*
- A PVC named *primary-pgdata*
- The database using predefined environment variables

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

To create the example and run the container:

```
cd $CCPROOT/examples/docker/primary
./run.sh
```

Connect from your local host as follows:

```
psql -h localhost -U testuser -W userdb
```

Kubernetes and OpenShift

To create the example:

```
cd $CCPROOT/examples/kube/primary
./run.sh
```

Connect from your local host as follows:

```
psql -h primary -U postgres postgres
```

Helm

This example resides under the `$CCPROOT/examples/helm` directory. View the README to run this example using Helm link:<https://github.com/CrunchyData/crunchy-containers/blob/master/examples/helm/primary/README.md>[here].

Statefulsets

This example deploys a statefulset named `statefulset`. The statefulset is a new feature in Kubernetes as of version 1.5 and in OpenShift Origin as of version 3.5. Statefulsets have replaced PetSets going forward.

Please view link:<https://kubernetes.io/docs/concepts/abstractions/controllers/statefulsets/>[this Kubernetes description] to better understand what a Statefulset is and how it works.

This example creates 2 PostgreSQL containers to form the set. At startup, each container will examine its hostname to determine if it is the first container within the set of containers.

The first container is determined by the hostname suffix assigned by Kubernetes to the pod. This is an ordinal value starting with `0`. If a container sees that it has an ordinal value of `0`, it will update the container labels to add a new label of:

```
name=$PG_PRIMARY_HOST
```

In this example, `PG_PRIMARY_HOST` is specified as `statefulset-primary`.

By default, the containers specify a value of `name=statefulset-replica`.

There are 2 services that end user applications will use to access the PostgreSQL cluster, one service (`statefulset-primary`) routes to the primary container and the other (`statefulset-replica`) to the replica containers.

```
$ ${CCP_CLI} get service
NAME                CLUSTER-IP      EXTERNAL-IP      PORT(S)          AGE
kubernetes          10.96.0.1       <none>           443/TCP          22h
statefulset-primary 10.97.168.138   <none>           5432/TCP         1h
statefulset-replica 10.97.218.221   <none>           5432/TCP         1h
```

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Kubernetes and OpenShift

First, start the example with the following command:

```
cd $CCPROOT/examples/kube/statefulset
./run.sh
```

You can access the primary database as follows:

```
psql -h statefulset-primary -U postgres postgres
```

You can access the replica databases as follows:

```
psql -h statefulset-replica -U postgres postgres
```

You can scale the number of containers using this command; this will essentially create an additional replica database.

```
`${CCP_CLI}` scale --replicas=3 statefulset statefulset
```

Helm

This example resides under the `$CCPROOT/examples/helm` directory. View the README to run this example using Helm link:<https://github.com/CrunchyData/crunchy-containers/blob/master/examples/helm/statefulset/README.md>[here].

Synchronous Replication

This example deploys a PostgreSQL cluster with a primary, a synchronous replica, and an asynchronous replica. The two replicas share the same service.

To shutdown the instance and remove the container for each example, run the following:

```
./cleanup.sh
```

Docker

To run this example, run the following:

```
cd $CCPROOT/examples/docker/sync
./run.sh
```

You can test the replication status on the primary by using the following command and the password `password`:

```
psql -h 127.0.0.1 -p 12010 -U postgres postgres -c 'table pg_stat_replication'
```

You should see 2 rows; 1 for the asynchronous replica and 1 for the synchronous replica. The `sync_state` column shows values of `async` or `sync`.

You can test replication to the replicas by first entering some data on the primary, and secondly querying the replicas for that data:

```
psql -h 127.0.0.1 -p 12010 -U postgres postgres -c 'create table foo (id int)'
psql -h 127.0.0.1 -p 12010 -U postgres postgres -c 'insert into foo values (1)'
psql -h 127.0.0.1 -p 12011 -U postgres postgres -c 'table foo'
psql -h 127.0.0.1 -p 12012 -U postgres postgres -c 'table foo'
```

Kubernetes and OpenShift

Running the example:

```
cd $CCPROOT/examples/kube/sync
./run.sh
```

Connect to the `primarysync` and `replicasync` databases as follows for both the Kubernetes and OpenShift environments:

```
psql -h primarysync -U postgres postgres -c 'create table test (id int)'
psql -h primarysync -U postgres postgres -c 'insert into test values (1)'
psql -h primarysync -U postgres postgres -c 'table pg_stat_replication'
psql -h replicasync -U postgres postgres -c 'select inet_server_addr(), * from test'
psql -h replicasync -U postgres postgres -c 'select inet_server_addr(), * from test'
psql -h replicasync -U postgres postgres -c 'select inet_server_addr(), * from test'
```

This set of queries will show you the IP address of the PostgreSQL replica container. Note the changing IP address due to the round-robin service proxy being used for both replicas. The example queries also show that both replicas are replicating successfully from the primary.

Requirements

These installation instructions are developed and tested for the following operating systems:

- **CentOS 7**
- **RHEL 7**

The Crunchy Container Suite can run on different environments including:

- **Docker 1.13+**
- **OpenShift Container Platform 3.11**
- **Kubernetes 1.8+**

In this document we list the basic installation steps required for these environments.

Project Environment

If your goal is to simply run the containers any properly configured user account should work. If your goal is for development and/or building the containers, we recommend a user whose environment is dedicated for that purpose.

First add the following lines to your `.bashrc` file to set the project paths:

```
export GOPATH=$HOME/cdev
export GOBIN=$GOPATH/bin
export PATH=$PATH:$GOBIN
export CCP_BASEOS=centos7          # centos7 for Centos, rhel7 for Redhat
export CCP_PGVERSION=10
export CCP_PG_FULLVERSION=10.6
export CCP_VERSION=2.2.0
export CCP_IMAGE_PREFIX=crunchydata
export CCP_IMAGE_TAG=$CCP_BASEOS-$CCP_PG_FULLVERSION-$CCP_VERSION
export CCROOT=$GOPATH/src/github.com/crunchydata/crunchy-containers
export CCP_SECURITY_CONTEXT=""
export CCP_CLI=kubectl            # kubectl for K8s, oc for OpenShift
export CCP_NAMESPACE=demo
```

You will need to add environment variables for storage configuration as well. Please see the [Storage Configuration](#) document for configuring storage using environment variables set in `.bashrc`.

It will be necessary to refresh your `.bashrc` file in order for the changes to take effect.

```
. ~/.bashrc
```

Next, set up a project directory structure and pull down the project from github:

```
mkdir -p $HOME/cdev/src/github.com/crunchydata $HOME/cdev/pkg $HOME/cdev/bin
```

Installation

The installation for CentOS 7 and RHEL 7 are similar, but there are several steps which require slightly different commands related to location of repositories, etc. These are highlighted below where necessary.

Install Supporting Software

CentOS 7 only

```
sudo yum -y install golang git docker
```

RHEL 7 only

```
sudo subscription-manager repos --enable=rhel-7-server-optional-rpms
sudo yum-config-manager --enable rhel-7-server-extras-rpms
sudo yum -y install git golang
```

Clone GitHub repository

```
cd $GOPATH/src/github.com/crunchydata
git clone https://github.com/crunchydata/crunchy-containers
cd crunchy-containers
git checkout 2.2.0
go get github.com/blang/expvar
```

If you are a Crunchy Enterprise Customer running on RHEL, you will place the Crunchy repository key and yum repository file into the `$CCPROOT/conf` directory at this point. These files can be obtained through <https://access.crunchydata.com/> on the downloads page.

Install PostgreSQL

These installation instructions assume the installation of PostgreSQL 10 through the official PGDG repository. View the documentation located [here](#) in order to view more detailed notes or install a different version of PostgreSQL.

Locate and edit your distribution's `.repo` file, located:

- On **CentOS**: `/etc/yum.repos.d/CentOS-Base.repo`, `[base]` and `[updates]` sections
- On **RHEL**: `/etc/yum/pluginconf.d/rhnplugin.conf` `[main]` section

To the section(s) identified above, depending on OS being used, you need to append a line to prevent dependencies from getting resolved to the PostgreSQL supplied by the base repository:

```
exclude=postgresql*
```

Next, install the RPM relating to the base operating system and PostgreSQL version you wish to install. The RPMs can be found [here](#). Below we chose Postgresql 10 for the example (change if you need different version):

On **CentOS** system:

```
sudo yum -y install
https://download.postgresql.org/pub/repos/yum/10/redhat/rhel-7-x86_64/pgdg-centos10-10-2.noarch.rpm
```

On **RHEL** system:

```
sudo yum -y install
https://download.postgresql.org/pub/repos/yum/testing/10/redhat/rhel-7-x86_64/pgdg-redhat10-10-2.n
```

Update the system:

```
sudo yum -y update
```

Install the PostgreSQL server package.

```
sudo yum -y install postgresql10-server.x86_64
```

Update the system:

```
sudo yum -y update
```

Install Docker

The OpenShift and Kubernetes (KubeAdm) instructions both have a section for installing docker. Installing docker now won't cause any issues but you may wish to configure Docker storage before bringing everything up. Configuring Docker Storage is different from *Storage Configuration* referenced earlier in the instructions and is not covered here.

For a basic docker installation, you can follow the instructions below. Please refer to the respective installation guide for the version of Kubernetes you are installing for more specific details.

```
sudo yum -y install docker
```

It is necessary to add the `docker` group and give your user access to that group:

```
sudo groupadd docker
sudo usermod -a -G docker <username>
```

Logout and login again as the same user to allow group settings to take effect.

Enable Docker service and start Docker (once all configuration is complete):

```
sudo systemctl enable docker.service
sudo systemctl start docker.service
```

OpenShift Installation

Use the OpenShift installation guide to install OpenShift Enterprise on your host. Make sure to choose the proper version of OpenShift you want to install. The main instructions for 3.11 are here and you'll be able to select a different version there, if needed:

<https://docs.openshift.com/container-platform/3.11/install/index.html>

Kubernetes Installation

Make sure your hostname resolves to a single IP address in your `/etc/hosts` file. The NFS examples will not work otherwise and other problems with installation can occur unless you have a resolving hostname.

You should see a single IP address returned from this command:

```
$ hostname --ip-address
```

Installing Kubernetes

We suggest using Kubeadm as a simple way to install Kubernetes.

See [Kubeadm](#) for installing the latest version of Kubeadm.

See [Create a Cluster](#) for creating the Kubernetes cluster using **Kubeadm**. Note: We find that Weave networking works particularly well with the container suite.

Please see [here](#) to view the official documentation regarding configuring DNS for your Kubernetes cluster.

Post Kubernetes Configuration

In order to run the various examples, Role Based Account Control will need to be set up. Specifically, the **cluster-admin** role will need to be assigned to the Kubernetes user that will be utilized to run the examples. This is done by creating the proper **ClusterRoleBinding**:

```
$ kubectl create clusterrolebinding cluster-admin-binding \
--clusterrole cluster-admin --user someuser
```

If you are running on GKE, the following command can be utilized to auto-populate the **user** option with the account that is currently logged into Google Cloud:

```
$ kubectl create clusterrolebinding cluster-admin-binding \
--clusterrole cluster-admin --user $(gcloud config get-value account)
```

If more than one user will be running the examples on the same Kubernetes cluster, a unique name will need to be provided for each new **ClusterRoleBinding** created in order to assign the **cluster-admin** role to every user. The example below will create a **ClusterRoleBinding** with a unique value:

```
$ kubectl create clusterrolebinding <unique>-cluster-admin-binding \  
--clusterrole cluster-admin \  
--user someuser
```

If you are running on GKE, the following can be utilized to create a unique **ClusterRoleBinding** for each user, with the user's Google Cloud account prepended to the name of each new **ClusterRoleBinding**:

```
$ kubectl create clusterrolebinding "$(gcloud config get-value account)-cluster-admin-binding" \  
--clusterrole cluster-admin \  
--user $(gcloud config get-value account)
```

Helm

Some Kubernetes Helm examples are provided in the following directory as one option for deploying the Container Suite.

```
$CCPROOT/examples/helm/
```

Once you have your Kubernetes environment configured, it is simple to get Helm up and running. Please refer to [this document](#) to get Helm installed and configured properly.

Configuring Namespace and Permissions

In Kubernetes, a concept called a **namespace** provides the means to separate created resources or components into individual logically grouped partitions. In OpenShift, *namespace* is referred to as a *project*.

It is considered a best practice to have dedicated namespaces for projects in both testing and production environments.

All examples in the Crunchy Container Suite operate within the namespace defined by the environment variable `$CCP_NAMESPACE`. The default we use for namespace is 'demo' but it can be set to any valid namespace name. The instructions below illustrate how to set up and work within new namespaces or projects in both Kubernetes and OpenShift.

Kubernetes

This section will illustrate how to set up a new Kubernetes namespace called **demo**, and will then show how to provide permissions to that namespace to allow the Kubernetes examples to run within that namespace.

First, view currently existing namespaces:

```
$ kubectl get namespace  
NAME          STATUS    AGE  
default       Active   21d  
kube-public   Active   21d  
kube-system   Active   21d
```

Then, create a new namespace called **demo**:

```
$ kubectl create -f $CCPROOT/conf/demo-namespace.json  
namespace "demo" created  
$ kubectl get namespace demo  
NAME          STATUS    AGE  
demo          Active    7s
```

Then set the namespace as the default for the current context:

When a namespace is not explicitly stated for a command, Kubernetes uses the namespace specified by the currently set context.

```
$ kubectl config set-context $(kubectl config current-context) --namespace=demo
```

We can verify that the namespace was set correctly through the following command:

```
$ kubectl config view | grep namespace:  
namespace: demo
```

OpenShift

This section assumes an administrator has already logged in first as the **system:admin** user as directed by the OpenShift Installation Guide.

For our development purposes only, we typically specify the OCP Authorization policy of **AllowAll** as documented here:

https://docs.openshift.com/container-platform/3.11/install_config/configuring_authentication.html#AllowAllPasswordIdentityProvider

We do not recommend this authentication policy for a production deployment of OCP.

For the best results, it is recommended that you run the examples with a user that has **NOT** been assigned the **cluster-admin** cluster role.

Log into the system as a user:

```
$ oc login -u <user>
```

The next step is to create a **demo** namespace to run the examples within. The name of this OCP project will be what you supply in the `CCP_NAMESPACE` environment variable:

```
$ oc new-project demo --description="Crunchy Containers project"
  --display-name="Crunchy-Containers"
Now using project "demo" on server "https://127.0.0.1:8443".
$ export CCP_NAMESPACE=demo
```

If we view the list of projects, we can see the new project has been added and is “active”.

```
$ oc get projects
NAME          DISPLAY NAME      STATUS
demo          Crunchy-Containers Active
myproject     My Project        Active
```

If you were on a different project and wanted to switch to the demo project, you would do so by running the following:

```
$ oc project demo
Now using project "demo" on server "https://127.0.0.1:8443".
```

When self-provisioning a new project using the `oc new-project` command, the current user (i.e., the user you used when logging into OCP with the `oc login` command) will automatically be assigned to the **admin** role for that project. This will allow the user to create the majority of the objects needed to successfully run the examples. However, in order to create the **Persistent Volume** objects needed to properly configure storage for the examples, an additional role is needed. Specifically, a new role is needed that can both create and delete **Persistent Volumes**.

Using the following two commands, create a new Cluster Role that has the ability to create and delete persistent volumes, and then assign that role to your current user:

Please be aware that the following two commands require privileges that your current user may not have. In the event that you are unable to run these commands, and do not have access to a user that is able to run them (e.g., the **system:admin** user that is created by default when installing OCP), please contact your local OCP administrator to run the commands on your behalf, or grant you the access required to run them yourself.

```
$ oc create clusterrole crunchytester --verb="list,create,delete" --resource=persistentvolumes
clusterrole "crunchytester" created

$ oc adm policy add-cluster-role-to-user crunchytester someuser
cluster role "crunchytester" added: "someuser"
```

Your user should now have the roles and privileges required to run the examples.

Storage Configuration

Available Storage Types

The Crunchy Container Suite is officially tested using two different storage backends:

- HostPath (single node testing)
- NFS (single and multi-node testing)

Other storage backends work as well, including GCE, EBS, ScaleIO, and others, but may require you to modify various examples or configuration.

The Crunchy Container Suite is tested, developed, and examples are provided that use the various storage types listed above. This ensures that customers have a high degree of choices when it comes to choosing a volume type. HostPath and NFS allow precise host path choices for where database volumes are persisted. HostPath and NFS also allow governance models where volume creation is performed by an administrator instead of the application/developer team.

Where customers desire a dynamic form of volume creation (e.g. self service), storage classes are also supported within the example set.

Environment variables are set to determine how and what storage is to be used.

NOTE: When running the examples using HostPath or NFS storage, the run scripts provided in the examples will create directories using the following pattern:

```
$CCP_STORAGE_PATH/$CCP_NAMESPACE-<EXAMPLE_NAME>
```

HostPath

HostPath is the simplest storage backend to setup. It is only feasible on a single node but is sufficient for testing the examples. In your `.bashrc` file, add the following variables to specify the proper settings for your the HostPath storage volume:

```
export CCP_SECURITY_CONTEXT=""
export CCP_STORAGE_PATH=/data
export CCP_STORAGE_MODE=ReadWriteMany
export CCP_STORAGE_CAPACITY=400M
```

NOTE: It may be necessary to grant your user in OpenShift or Kubernetes the rights to modify the `hostaccess` SCC. This can be done with the following command:

```
oadm policy add-scc-to-user hostaccess $(oc whoami)
```

NFS

NFS can also be utilized as a storage mechanism. Instructions for setting up a NFS can be found in the **Configuration Notes for NFS** section below.

For testing with NFS, include the following variables in your `.bashrc` file, providing the proper configuration details for your NFS:

```
export CCP_SECURITY_CONTEXT='"supplementalGroups": [65534] '
export CCP_STORAGE_PATH=/nfsfileshare
export CCP_NFS_IP=<IP OF NFS SERVER>
export CCP_STORAGE_MODE=ReadWriteMany
export CCP_STORAGE_CAPACITY=400M
```

In the example above the group ownership of the NFS mount is assumed to be `nfsnobody` or `65534`. Additionally, it is recommended that root not be squashed on the NFS share (using `no_root_squash`) in order to ensure the proper directories can be created, modified and removed as needed for the various container examples.

Additionally, the examples in the Crunchy Container suite need access to the NFS in order to create the directories utilized by the examples. The NFS should therefore be mounted locally so that the `run.sh` scripts contained within the examples can complete the proper setup.

Configuration Notes for NFS

- Most of the Crunchy containers run as the postgres UID (26), but you will notice that when `supplementalGroups` is specified, the pod will include the `nfsnobody` group in the list of groups for the pod user
- If you are running your NFS system with SELinux in enforcing mode, you will need to run the following command to allow NFS write permissions:

```
sudo setsebool -P virt_use_nfs 1
```

- Detailed instructions for setting up a NFS server on Centos 7 can be found using the following link:
<http://www.itzgeek.com/how-tos/linux/centos-how-tos/how-to-setup-nfs-server-on-centos-7-rhel-7-fedora-22.html>
- If you are running your client on a VM, you will need to add `insecure` to the `exports` file on the NFS server due to the way port translation is done between the VM host and the VM instance. For more details on this bug, please see the following link:
<http://serverfault.com/questions/107546/mount-nfs-access-denied-by-server-while-mounting>

- A suggested best practice for tuning NFS for PostgreSQL is to configure the PostgreSQL fstab mount options like so:

```
proto=tcp,suid,rw,vers=3,proto=tcp,timeo=600,retrans=2,hard,fg,rsize=8192,wsiz=8192
```

And to then change your network options as follows:

```
MTU=9000
```

- If interested in mounting the same NFS share multiple times on the same mount point, look into the [noac mount option](#)

Dynamic Storage

Dynamic storage classes can be used for the examples. There are various providers and solutions for dynamic storage, so please consult the Kubernetes documentation for additional details regarding supported storage choices. The environment variable `CCP_STORAGE_CLASS` is used in the examples to determine whether or not to create a PersistentVolume manually, or if it will be created dynamically using a StorageClass. In the case of GKE, the default StorageClass is named **default**. Storage class names are determined by the Kubernetes administrator and can vary.

Using block storage requires a security context to be set as follows:

```
export CCP_SECURITY_CONTEXT='"fsGroup":26'  
export CCP_STORAGE_CLASS=standard  
export CCP_STORAGE_MODE=ReadWriteOnce  
export CCP_STORAGE_CAPACITY=400M
```