



Documentation

1.12.0 (December 7, 2022)

Percona Technical Documentation Team

Percona LLC and/or its affiliates, © 2009 - 2022

Table of contents

1. Percona Operator for MySQL based on Percona XtraDB Cluster	4
2. Requirements	4
3. Quickstart guides	4
4. Advanced Installation Guides	4
5. Configuration	4
6. Management	5
7. HOWTOs	5
8. Reference	5
9. Requirements	6
9.1 System Requirements	6
9.2 Design overview	7
9.3 Compare various solutions to deploy MySQL in Kubernetes	9
10. Quickstart guides	12
10.1 Install Percona XtraDB Cluster using Helm	12
10.2 Install Percona XtraDB Cluster on Minikube	13
11. Advanced Installation guide	18
11.1 Install Percona XtraDB Cluster on Google Kubernetes Engine (GKE)	18
11.2 Install Percona XtraDB Cluster on Amazon Elastic Kubernetes Service (EKS)	22
11.3 Install Percona XtraDB Cluster on Azure Kubernetes Service (AKS)	25
11.4 Install the Operator and deploy your Percona XtraDB Cluster	25
11.5 Install Percona XtraDB Cluster on Kubernetes	29
11.6 Install Percona XtraDB Cluster on OpenShift	33
11.7 Use docker images from a custom registry	37
12. Configuration	40
12.1 Users	40
12.2 Changing MySQL Options	43
12.3 Binding Percona XtraDB Cluster components to Specific Kubernetes/OpenShift Nodes	46
12.4 Labels and annotations	49
12.5 Local Storage support for the Percona Operator for MySQL	50
12.6 . Now you can deploy the Operator and Percona XtraDB Cluster using this	51
12.7 Define environment variables	52
12.8 Configuring Load Balancing with HAProxy	54
12.9 Configuring Load Balancing with ProxySQL	58
12.10 Transport Layer Security (TLS)	63
12.11 Data at Rest Encryption	69

12.12 Telemetry	72
13. Management	73
13.1 Providing Backups	73
13.2 Upgrade Database and Operator	85
13.3 Scale Percona XtraDB Cluster on Kubernetes and OpenShift	89
13.4 Set up Percona XtraDB Cluster cross-site replication	91
13.5 Monitoring	94
13.6 Using sidecar containers	97
13.7 Pause/resume Percona XtraDB Cluster	100
13.8 Crash Recovery	101
13.9 Debug	105
14. HOWTOs	107
14.1 Install Percona XtraDB Cluster in multi-namespace (cluster-wide) mode	107
14.2 How to upgrade Percona XtraDB Cluster manually	111
15. Reference	113
15.1 Custom Resource options	113
15.2 Percona certified images	151
15.3 Percona Operator for MySQL API Documentation	154
15.4 Frequently Asked Questions	194
15.5 Old releases (documentation archive)	198
16. Release Notes	199
16.1 <i>Percona Operator for MySQL based on Percona XtraDB Cluster 1.12.0</i>	199
16.2 <i>Percona Operator for MySQL based on Percona XtraDB Cluster 1.11.0</i>	201
16.3 <i>Percona Distribution for MySQL Operator 1.10.0</i>	203
16.4 <i>Percona Distribution for MySQL Operator 1.9.0</i>	205
16.5 <i>Percona Kubernetes Operator for Percona XtraDB Cluster 1.8.0</i>	207
16.6 <i>Percona Kubernetes Operator for Percona XtraDB Cluster 1.7.0</i>	209
16.7 <i>Percona Kubernetes Operator for Percona XtraDB Cluster 1.6.0</i>	211
16.8 <i>Percona Kubernetes Operator for Percona XtraDB Cluster 1.5.0</i>	213
16.9 <i>Percona Kubernetes Operator for Percona XtraDB Cluster 1.4.0</i>	215
16.10 <i>Percona Kubernetes Operator for Percona XtraDB Cluster 1.3.0</i>	216
16.11 <i>Percona Kubernetes Operator for Percona XtraDB Cluster 1.2.0</i>	217
16.12 <i>Percona Kubernetes Operator for Percona XtraDB Cluster 1.1.0</i>	218
16.13 <i>Percona Kubernetes Operator for Percona XtraDB Cluster 1.0.0</i>	219

1. Percona Operator for MySQL based on Percona XtraDB Cluster

Kubernetes and the OpenShift platform, based on Kubernetes, have added a way to manage containerized systems, including database clusters. This management is achieved by controllers, declared in configuration files. These controllers provide automation with the ability to create objects, such as a container or a group of containers called pods, to listen for an specific event and then perform a task.

This automation adds a level of complexity to the container-based architecture and stateful applications, such as a database. A Kubernetes Operator is a special type of controller introduced to simplify complex deployments. The Operator extends the Kubernetes API with custom resources.

[Percona XtraDB Cluster](#) is an open-source enterprise MySQL solution that helps you to ensure data availability for your applications while improving security and simplifying the development of new applications in the most demanding public, private, and hybrid cloud environments.

Following our best practices for deployment and configuration, [Percona Operator for MySQL based on Percona XtraDB Cluster](#) contains everything you need to quickly and consistently deploy and scale Percona XtraDB Cluster instances in a Kubernetes-based environment on-premises or in the cloud.

2. Requirements

- [System Requirements](#)
- [Design and architecture](#)
- [Comparison with other solutions](#)

3. Quickstart guides

- [Install with Helm](#)
- [Install on Minikube](#)
- [Install on Google Kubernetes Engine \(GKE\)](#)
- [Install on Amazon Elastic Kubernetes Service \(AWS EKS\)](#)
- [Install on Microsoft Azure Kubernetes Service \(AKS\)](#)

4. Advanced Installation Guides

- [Generic Kubernetes installation](#)
- [Install on OpenShift](#)
- [Use private registry](#)

5. Configuration

- [Application and system users](#)
- [Changing MySQL Options](#)
- [Anti-affinity and tolerations](#)
- [Labels and annotations](#)
- [Local Storage support](#)
- [Defining environment variables](#)

- [Load Balancing with HAProxy](#)
- [Load Balancing with ProxySQL](#)
- [Transport Encryption \(TLS/SSL\)](#)
- [Data at rest encryption](#)
- [Telemetry](#)

6. Management

- [Backup and restore](#)
- [Upgrade Database and Operator](#)
- [Horizontal and vertical scaling](#)
- [Multi-cluster and multi-region deployment](#)
- [Monitor with Percona Monitoring and Management \(PMM\)](#)
- [Add sidecar containers](#)
- [Restart or pause the cluster](#)
- [Crash recovery](#)
- [Debug and troubleshoot](#)

7. HOWTOS

- [How to install Percona XtraDB Cluster in multi-namespace \(cluster-wide\) mode](#)
- [How to upgrade Percona XtraDB Cluster manually](#)

8. Reference

- [Custom Resource options](#)
- [Percona certified images](#)
- [Operator API](#)
- [Frequently Asked Questions](#)
- [Old releases \(documentation archive\)](#)
- [Release Notes](#)

Last update: 2022-12-07

9. Requirements

9.1 System Requirements

The Operator supports Percona XtraDB Cluster (PXC) 5.7 and 8.0.

The new `caching_sha2_password` authentication plugin which is default in 8.0 is not supported for the ProxySQL compatibility reasons. Therefore both Percona XtraDB Cluster 5.7 and 8.0 are configured with `default_authentication_plugin = mysql_native_password`.

9.1.1 Officially supported platforms

The following platforms were tested and are officially supported by the Operator 1.12.0:

- [Google Kubernetes Engine \(GKE\)](#) 1.21 - 1.24
- [Amazon Elastic Container Service for Kubernetes \(EKS\)](#) 1.21 - 1.24
- [Azure Kubernetes Service \(AKS\)](#) 1.22 - 1.24
- [OpenShift](#) 4.10 - 4.11
- [Minikube](#) 1.28

Other Kubernetes platforms may also work but have not been tested.

9.1.2 Resource Limits

A cluster running an officially supported platform contains at least three Nodes, with the following resources:

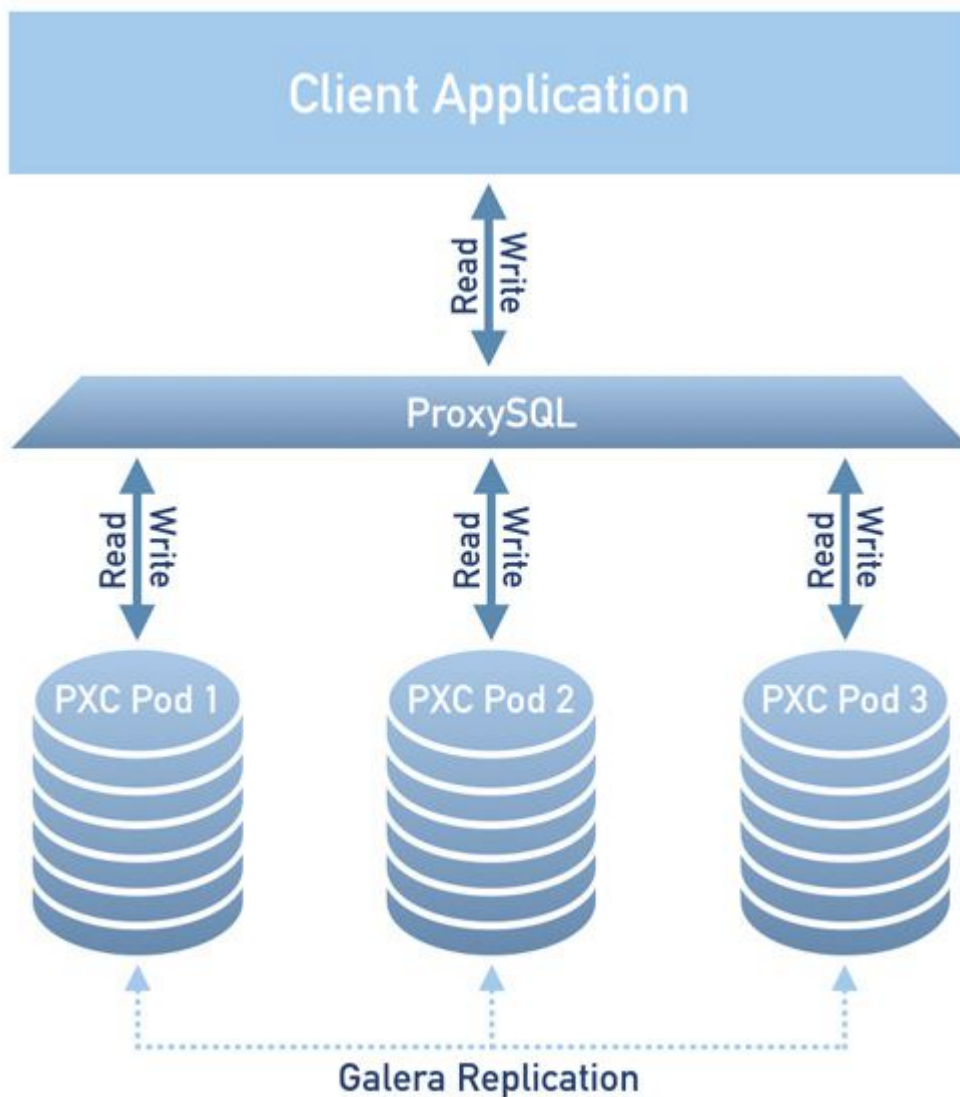
- 2GB of RAM,
- 2 CPU threads per Node for Pods provisioning,
- at least 60GB of available storage for Persistent Volumes provisioning.

Last update: 2022-11-25

9.2 Design overview

Percona XtraDB Cluster integrates *Percona Server for MySQL* running with the XtraDB storage engine, and *Percona XtraBackup* with the *Galera library* to enable synchronous multi-primary replication.

The design of the Operator is highly bound to the Percona XtraDB Cluster high availability implementation, which in its turn can be briefly described with the following diagram.




Being a regular MySQL Server instance, each node contains the same set of data synchronized across nodes. The recommended configuration is to have at least 3 nodes. In a basic setup with this amount of nodes, Percona XtraDB Cluster provides high availability, continuing to function if you take any of the nodes down. Additionally load balancing can be achieved with the ProxySQL daemon, which accepts incoming traffic from MySQL clients and forwards it to backend MySQL servers.

Note

Using ProxySQL results in [more efficient database workload management](#) in comparison with other load balancers which are not SQL-aware, including built-in ones of the cloud providers, or the Kubernetes NGINX Ingress Controller.

To provide high availability operator uses [node affinity](#) to run Percona XtraDB Cluster instances on separate worker nodes if possible. If some node fails, the pod with it is automatically re-created on another node.



To provide data storage for stateful applications, Kubernetes uses Persistent Volumes. A *PersistentVolumeClaim* (PVC) is used to implement the automatic storage provisioning to pods. If a failure occurs, the Container Storage Interface (CSI) should be able to re-mount storage on a different node. The PVC StorageClass must support this feature (Kubernetes and OpenShift support this in versions 1.9 and 3.9 respectively).

The Operator functionality extends the Kubernetes API with *PerconaXtraDBCluster* object, and it is implemented as a golang application. Each *PerconaXtraDBCluster* object maps to one separate Percona XtraDB Cluster setup. The Operator listens to all events on the created objects. When a new *PerconaXtraDBCluster* object is created, or an existing one undergoes some changes or deletion, the operator automatically creates/changes/deletes all needed Kubernetes objects with the appropriate settings to provide a proper Percona XtraDB Cluster operation.

Last update: 2022-09-22

9.3 Compare various solutions to deploy MySQL in Kubernetes

There are multiple ways to deploy and manage MySQL in Kubernetes. Here we will focus on comparing the following open source solutions:

- [Bitnami Helm chart](#)
- [KubeDB](#)
- [MongoDB Community Operator](#)
- [Percona Operator for MongoDB](#)
- [Bitpoke MySQL Operator \(former Presslabs\)](#)
- [Oracle MySQL Operator](#)
- [Moco](#) by Cybozu
- [Vitess Operator](#) by PlanetScale
- [Percona Operator for MySQL](#)
 - [based on Percona XtraDB Cluster](#)
 - [based on Percona Server for MySQL](#)

9.3.1 Generic

The review of generic features, such as supported MySQL versions, open source models and more.

Feature/ Product	Percona Operator for MySQL (based on PXC)	Percona Operator for MySQL (based on PS)	Bitpoke MySQL Operator	Moco	Oracle MySQL Operator	Vitess
Open source model	Apache 2.0	Apache 2.0	Apache 2.0	Apache 2.0	Apache 2.0	Apache 2.0
MySQL versions	5.7, 8.0	8.0	5.7	8.0	8.0	5.7, 8.0
Kubernetes conformance	Various versions are tested	Various versions are tested	Not guaranteed	Not guaranteed	Not guaranteed	Not guaranteed
Paid support	✓	✓	⊘	⊘	✓	⊘

9.3.2 MySQL Topologies

Focus on replication capabilities and proxies integrations.

Feature/Product	Percona Operator for MySQL (based on PXC)	Percona Operator for MySQL (based on PS)	Bitpoke MySQL Operator	Moco	Oracle MySQL Operator	Vitess
Replication	Sync with Galera	Async and Group Replication	Async	Semi-sync	Group Replication	Async
Proxy	HAProxy and ProxySQL	HAProxy and MySQL Router	None	None	MySQL Router	VTGate
Multi-cluster deployment	✓	⊘	⊘	⊘	⊘	⊘
Sharding	⊘	⊘	⊘	⊘	⊘	✓

9.3.3 Backups

Here are the backup and restore capabilities of each solution.

Feature/Product	Percona Operator for MySQL (based on PXC)	Percona Operator for MySQL (based on PS)	Bitpoke MySQL Operator	Moco	Oracle MySQL Operator	Vitess
Scheduled backups	✓	✓	✓	✓	⊘	✓
Incremental backups	⊘	⊘	⊘	✓	⊘	⊘
PITR	✓	✓	⊘	⊘	⊘	⊘
PVCs for backups	✓	⊘	⊘	⊘	⊘	⊘

9.3.4 Monitoring

Monitoring is crucial for any operations team.

Feature/Product	Percona Operator for MySQL (based on PXC)	Percona Operator for MySQL (based on PS)	Bitpoke MySQL Operator	Moco	Oracle MySQL Operator	Vitess
Custom exporters	Through sidecars	Through sidecars	mysqld_exporter	mysqld_exporter	⊘	⊘
PMM	✓	✓	⊘	⊘	⊘	⊘

9.3.5 Miscellaneous

Compare various features that are not a good fit for other categories.

Feature/ Product	Percona Operator for MySQL (based on PXC)	Percona Operator for MySQL (based on PS)	Bitpoke MySQL Operator	Moco	Oracle MySQL Operator	Vitess
Customize MySQL	ConfigMaps and Secrets	ConfigMaps and Secrets	ConfigMaps	ConfigMaps	ConfigMaps	⊘
Helm	✓	✓	✓	✓	✓	⊘
Transport encryption	✓	✓	⊘	⊘	✓	✓
Encryption- at-rest	✓	✓	⊘	⊘	⊘	⊘

Last update: 2022-11-03

10. Quickstart guides

10.1 Install Percona XtraDB Cluster using Helm

Helm is the package manager for Kubernetes. Percona Helm charts can be found in [percona/percona-helm-charts](https://github.com/percona/percona-helm-charts) repository on Github.

10.1.1 Pre-requisites

Install Helm following its [official installation instructions](#).

Note

Helm v3 is needed to run the following steps.

10.1.2 Installation

1. Add the Percona's Helm charts repository and make your Helm client up to date with it:

```
$ helm repo add percona https://percona.github.io/percona-helm-charts/
$ helm repo update
```

2. Install the Percona Operator for MySQL based on Percona XtraDB Cluster:

```
$ helm install my-op percona/pxc-operator
```

The `my-op` parameter in the above example is the name of a [new release object](#) which is created for the Operator when you install its Helm chart (use any name you like).

Note

If nothing explicitly specified, `helm install` command will work with `default` namespace. To use different namespace, provide it with the following additional parameter: `--namespace my-namespace`.

3. Install Percona XtraDB Cluster:

```
$ helm install my-db percona/pxc-db
```

The `my-db` parameter in the above example is the name of a [new release object](#) which is created for the Percona XtraDB Cluster when you install its Helm chart (use any name you like).

10.1.3 Installing Percona XtraDB Cluster with customized parameters

The command above installs Percona XtraDB Cluster with [default parameters](#). Custom options can be passed to a `helm install` command as a `--set key=value[,key=value]` argument. The options passed with a chart can be any of the Operator's [Custom Resource options](#).

The following example will deploy a Percona XtraDB Cluster in the `pxc` namespace, with disabled backups and 20 Gi storage:

```
$ helm install my-db percona/pxc-db \
--set pxc.volumeSpec.resources.requests.storage=20Gi \
--set backup.enabled=false
```

Last update: 2022-12-07

10.2 Install Percona XtraDB Cluster on Minikube

Installing the Percona Operator for MySQL based on Percona XtraDB Cluster on [minikube](#) is the easiest way to try it locally without a cloud provider. Minikube runs Kubernetes on GNU/Linux, Windows, or macOS system using a system-wide hypervisor, such as VirtualBox, KVM/QEMU, VMware Fusion or Hyper-V. Using it is a popular way to test the Kubernetes application locally prior to deploying it on a cloud.

The following steps are needed to run the Operator and Percona XtraDB Cluster on Minikube:

1. [Install Minikube](#), using a way recommended for your system. This includes the installation of the following three components:
 - a. kubectl tool,
 - b. a hypervisor, if it is not already installed,
 - c. actual Minikube package.

After the installation, run `minikube start --memory=4096 --cpus=3` (parameters increase the virtual machine limits for the CPU cores and memory, to ensure stable work of the Operator). Being executed, this command will download needed virtualized images, then initialize and run the cluster. After Minikube is successfully started, you can optionally run the Kubernetes dashboard, which visually represents the state of your cluster. Executing `minikube dashboard` will start the dashboard and open it in your default web browser.

2. Deploy the operator with the following command:

```
$ kubectl apply -f https://raw.githubusercontent.com/percona/percona-xtradb-cluster-operator/v1.12.0/deploy/bundle.yaml
```

3. Deploy Percona XtraDB Cluster:

```
$ kubectl apply -f https://raw.githubusercontent.com/percona/percona-xtradb-cluster-operator/v1.12.0/deploy/cr-minimal.yaml
```

Note

This deploys one Percona XtraDB Cluster node and one HAProxy node. The [deploy/cr-minimal.yaml](#) is for minimal non-production deployment. For more configuration options please see [deploy/cr.yaml](#) and [Custom Resource Options](#). You can clone the repository with all manifests and source code by executing the following command:

```
$ git clone -b v1.12.0 https://github.com/percona/percona-xtradb-cluster-operator
```

After editing the needed options, apply your modified `deploy/cr.yaml` file as follows:

```
$ kubectl apply -f deploy/cr.yaml
```

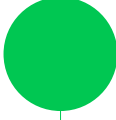
Creation process will take some time. The process is over when both operator and replica set pod have reached their Running status. `kubectl get pods` output should look like this:

NAME	READY	STATUS	RESTARTS	AGE
percona-xtradb-cluster-operator-d99c748-sqddq	1/1	Running	0	49m
minimal-cluster-pxc-0	3/3	Running	0	47m
minimal-cluster-haproxy-0	2/2	Running	0	47m

You can also track the progress via the Kubernetes dashboard:

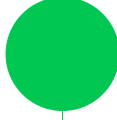
Workloads

Workload Status



Running: 1

Deployments



Running: 3

Pods



Running: 1

Replica Sets



Running: 2

Stateful Sets

Deployments

Name	Images	Labels	Pods	Created ↑
percona-xtradb-cluster-operator	percona/percona-xtradb-cluster-operator:1.11.0	-	1 / 1	4 minutes ago

Pods

Name	Images	Labels	Node	Status	Restarts	CPU Usage (cores)	Memory Usage (bytes)	Created ↑
minimal-cluster-haproxy-0	percona/percona-xtradb-cluster-operator:1.11.0-haproxy	app.kubernetes.io/component: haproxy	minikube	Running	0	-	-	4 minutes ago
	percona/percona-xtradb-cluster-operator:1.11.0-haproxy	app.kubernetes.io/instance: minimal-cluster						
	percona/percona-xtradb-cluster-operator:1.11.0-haproxy	app.kubernetes.io/managed-by: percona-xtradb-cluster-operator						
	percona/percona-xtradb-cluster-operator:1.11.0-logcollector	app.kubernetes.io/component: pxc						

4. During previous steps, the Operator has generated several [secrets](#), including the password for the `root` user, which you will definitely need to access the cluster. Use `kubectl get secrets` to see the list of Secrets objects (by default Secrets object you are interested in has `minimal-cluster-secrets` name). Then `kubectl get secret minimal-cluster-secrets -o yaml` will return the YAML file with generated secrets, including the root password which should look as follows:

```
...
data:
  ...
  root: cm9vdF9wYXNzd29yZA==
```

Here the actual password is base64-encoded, and `echo 'cm9vdF9wYXNzd29yZA==' | base64 --decode` will bring it back to a human-readable form.

5. Check connectivity to a newly created cluster.

First of all, run a MySQL client container and connect its console output to your terminal (running it may require some time to deploy the correspondent Pod). The following command will do this, naming the new Pod `percona-client`:

```
$ kubectl run -i --rm --tty percona-client --image=percona:8.0 --restart=Never -- bash -il
```

Now run `mysql` tool in the `percona-client` command shell using the password obtained from the secret:

```
$ mysql -h minimal-cluster-haproxy -uroot -proot_password
```

This command will connect you to the MySQL monitor.

```
mysql: [Warning] Using a password on the command line interface can be insecure.
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1872
Server version: 8.0.22-13.1 Percona XtraDB Cluster (GPL), Release rel13, Revision a48e6d5, WSREP version 26.4.3

Copyright (c) 2009-2021 Percona LLC and/or its affiliates
Copyright (c) 2000, 2021, Oracle and/or its affiliates.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
```


Last update: 2022-12-04

11. Advanced Installation guide

11.1 Install Percona XtraDB Cluster on Google Kubernetes Engine (GKE)

This quickstart shows you how to configure the Percona Operator for MySQL based on Percona XtraDB Cluster with the Google Kubernetes Engine. The document assumes some experience with Google Kubernetes Engine (GKE). For more information on the GKE, see the [Kubernetes Engine Quickstart](#).

11.1.1 Prerequisites

All commands from this quickstart can be run either in the **Google Cloud shell** or in **your local shell**.

To use *Google Cloud shell*, you need nothing but a modern web browser.

If you would like to use *your local shell*, install the following:

1. [gcloud](#). This tool is part of the Google Cloud SDK. To install it, select your operating system on the [official Google Cloud SDK documentation page](#) and then follow the instructions.
2. [kubectl](#). It is the Kubernetes command-line tool you will use to manage and deploy applications. To install the tool, run the following command:

```
$ gcloud auth login
$ gcloud components install kubectl
```

11.1.2 Configuring default settings for the cluster

You can configure the settings using the `gcloud` tool. You can run it either in the [Cloud Shell](#) or in your local shell (if you have installed Google Cloud SDK locally on the previous step). The following command will create a cluster named `my-cluster-1`:

```
$ gcloud container clusters create my-cluster-1 --project <project name> --zone us-central1-a --cluster-version 1.23 --machine-type n1-standard-4 --num-nodes=3
```

Note

You must edit the following command and other command-line statements to replace the `<project name>` placeholder with your project name. You may also be required to edit the *zone location*, which is set to `us-central1` in the above example. Other parameters specify that we are creating a cluster with 3 nodes and with machine type of 4 vCPUs and 45 GB memory.

You may wait a few minutes for the cluster to be generated, and then you will see it listed in the Google Cloud console (select *Kubernetes Engine* → *Clusters* in the left menu panel):

<input type="checkbox"/>	<input checked="" type="checkbox"/>	my-cluster-1	us-central1-a	3	12 vCPUs	45.00 GB	Connect		
--------------------------	-------------------------------------	--------------	---------------	---	----------	----------	---------	--	--

Now you should configure the command-line access to your newly created cluster to make `kubectl` be able to use it.

In the Google Cloud Console, select your cluster and then click the *Connect* shown on the above image. You will see the connect statement configures command-line access. After you have edited the statement, you may run the command in your local shell:

```
$ gcloud container clusters get-credentials my-cluster-1 --zone us-central1-a --project <project name>
```

11.1.3 Installing the Operator

1. First of all, use your [Cloud Identity and Access Management \(Cloud IAM\)](#) to control access to the cluster. The following command will give you the ability to create Roles and RoleBindings:

```
$ kubectl create clusterrolebinding cluster-admin-binding --clusterrole cluster-admin --user $(gcloud config get-value core/account)
```

The return statement confirms the creation:

```
clusterrolebinding.rbac.authorization.k8s.io/cluster-admin-binding created
```

2. Create a namespace and set the context for the namespace. The resource names must be unique within the namespace and provide a way to divide cluster resources between users spread across multiple projects.

So, create the namespace and save it in the namespace context for subsequent commands as follows (replace the `<namespace name>` placeholder with some descriptive name):

```
$ kubectl create namespace <namespace name>
$ kubectl config set-context $(kubectl config current-context) --namespace=<namespace name>
```

At success, you will see the message that namespace/ was created, and the context (gke_) was modified.

3. Use the following `git clone` command to download the correct branch of the `percona-xtradb-cluster-operator` repository:

```
$ git clone -b v1.12.0 https://github.com/percona/percona-xtradb-cluster-operator
```

After the repository is downloaded, change the directory to run the rest of the commands in this document:

```
$ cd percona-xtradb-cluster-operator
```

4. Deploy the Operator with the following command:

```
$ kubectl apply -f deploy/bundle.yaml
```

The following confirmation is returned:

```
customresourcedefinition.apiextensions.k8s.io/perconaxtradbclusters.pxc.percona.com created
customresourcedefinition.apiextensions.k8s.io/perconaxtradbclusterbackups.pxc.percona.com created
customresourcedefinition.apiextensions.k8s.io/perconaxtradbclusterrestores.pxc.percona.com created
customresourcedefinition.apiextensions.k8s.io/perconaxtradbbackups.pxc.percona.com created
role.rbac.authorization.k8s.io/percona-xtradb-cluster-operator created
serviceaccount/percona-xtradb-cluster-operator created
rolebinding.rbac.authorization.k8s.io/service-account-percona-xtradb-cluster-operator created
deployment.apps/percona-xtradb-cluster-operator created
```

5. The operator has been started, and you can create the Percona XtraDB cluster:

```
$ kubectl apply -f deploy/cr.yaml
```

The process could take some time. The return statement confirms the creation:

```
perconaxtradbcluster.pxc.percona.com/cluster1 created
```

6. During previous steps, the Operator has generated several [secrets](#), including the password for the `root` user, which you will need to access the cluster.

Use `kubectl get secrets` command to see the list of Secrets objects (by default Secrets object you are interested in has `cluster1-secrets` name). Then `kubectl get secret cluster1-secrets -o yaml` will return the YAML file with generated secrets, including the root password which should look as follows:

```
...
data:
  ...
  root: cm9vdF9wYXNzd29yZA==
```

Here the actual password is base64-encoded, and `echo 'cm9vdF9wYXNzd29yZA==' | base64 --decode` will bring it back to a human-readable form.

11.1.4 Verifying the cluster operator

It may take ten minutes to get the cluster started. You can verify its creation with the `kubectl get pods` command:

NAME	READY	STATUS	RESTARTS	AGE
cluster1-haproxy-0	2/2	Running	0	6m17s
cluster1-haproxy-1	2/2	Running	0	4m59s
cluster1-haproxy-2	2/2	Running	0	4m36s

```

cluster1-pxc-0          3/3   Running  0         6m17s
cluster1-pxc-1          3/3   Running  0         5m3s
cluster1-pxc-2          3/3   Running  0         3m56s
percona-xtradb-cluster-operator-79966668bd-rswbk 1/1   Running  0         9m54s

```

Also, you can see the same information when browsing Pods of your cluster in Google Cloud console via the *Object Browser*:

Name	Status	Type	Cluster
▼ core			API Group
▼ Pod			Kind
cluster1-haproxy-0	✔ Running	Pod	my-cluster-1
cluster1-haproxy-1	✔ Running	Pod	my-cluster-1
cluster1-haproxy-2	✔ Running	Pod	my-cluster-1
cluster1-pxc-0	✔ Running	Pod	my-cluster-1
cluster1-pxc-1	✔ Running	Pod	my-cluster-1
cluster1-pxc-2	✔ Running	Pod	my-cluster-1

If all nodes are up and running, you can try to connect to the cluster. Run a container with `mysql` tool and connect its console output to your terminal. The following command will do this, naming the new Pod `percona-client`:

```
$ kubectl run -i --rm --tty percona-client --image=percona:8.0 --restart=Never -- bash -il
```

Executing this command will open a `bash` command prompt:

```

If you don't see a command prompt, try pressing enter.
$

```

Now run `mysql` tool in the `percona-client` command shell using the password obtained from the secret:

```
$ mysql -h cluster1-haproxy -uroot -proot_password
```

This command will connect you to the MySQL monitor.

```

mysql: [Warning] Using a password on the command line interface can be insecure.
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1976
Server version: 8.0.19-10 Percona XtraDB Cluster (GPL), Release rel10, Revision 727f180, WSREP version 26.4.3

Copyright (c) 2009-2020 Percona LLC and/or its affiliates
Copyright (c) 2000, 2020, Oracle and/or its affiliates. All rights reserved.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

```

The following example will use the MySQL prompt to check the `max_connections` variable:

```
mysql> SHOW VARIABLES LIKE "max_connections";
```

The return statement displays the current `max_connections`.

```

+-----+-----+
| Variable_name | Value |
+-----+-----+
| max_connections | 79 |
+-----+-----+
1 row in set (0.02 sec)

```

11.1.5 Troubleshooting

If `kubectl get pods` command had shown some errors, you can examine the problematic Pod with the `kubectl describe <pod name>` command. For example, this command returns information for the selected Pod:

```
$ kubectl describe pod cluster1-haproxy-2
```

Review the detailed information for `Warning` statements and then correct the configuration. An example of a warning is as follows:

Warning FailedScheduling 68s (x4 over 2m22s) default-scheduler 0/1 nodes are available: 1 node(s) didn't match pod affinity/anti-affinity, 1 node(s) didn't satisfy existing pods anti-affinity rules.

Alternatively, you can examine your Pods via the *object browser*. Errors will look as follows:

Name	Status	Type	Cluster
▼ core		API Group	
▼ Pod		Kind	
cluster1-haproxy-0	✔ Running	Pod	my-cluster-1
cluster1-haproxy-1	✔ Running	Pod	my-cluster-1
cluster1-haproxy-2	❗ Unschedulable	Pod	my-cluster-1
cluster1-pxc-0	✔ Running	Pod	my-cluster-1
cluster1-pxc-1	✔ Running	Pod	my-cluster-1
cluster1-pxc-2	❗ Unschedulable	Pod	my-cluster-1

Clicking the problematic Pod will bring you to the details page with the same warning:

❗ cluster1-haproxy-2

❗ 0/2 nodes are available: 2 node(s) didn't match pod affinity/anti-affinity, 2 node(s) didn't satisfy existing pods anti-affinity rules. Show Details

[Details](#) [Events](#) [Logs](#) [YAML](#)

1h 6h 1d 7d 30d

11.1.6 Removing the GKE cluster

There are several ways that you can delete the cluster.

You can clean up the cluster with the `gcloud` command as follows:

```
$ gcloud container clusters delete <cluster name>
```

The return statement requests your confirmation of the deletion. Type `y` to confirm.

Also, you can delete your cluster via the GKE console. Just click the appropriate trashcan icon in the clusters list:

<input type="checkbox"/>	✔ my-cluster-1	us-central1-a	3	12 vCPUs	45.00 GB	Connect	
--------------------------	----------------	---------------	---	----------	----------	---------	--

The cluster deletion may take time.

Last update: 2022-12-04

11.2 Install Percona XtraDB Cluster on Amazon Elastic Kubernetes Service (EKS)

This quickstart shows you how to deploy the Operator and Percona XtraDB Cluster on Amazon Elastic Kubernetes Service (EKS). The document assumes some experience with Amazon EKS. For more information on the EKS, see the [Amazon EKS official documentation](#).

11.2.1 Prerequisites

The following tools are used in this guide and therefore should be preinstalled:

1. **AWS Command Line Interface (AWS CLI)** for interacting with the different parts of AWS. You can install it following the [official installation instructions for your system](#).
2. **eksctl** to simplify cluster creation on EKS. It can be installed along its [installation notes on GitHub](#).
3. **kubectl** to manage and deploy applications on Kubernetes. Install it [following the official installation instructions](#).

Also, you need to configure AWS CLI with your credentials according to the [official guide](#).

11.2.2 Create the EKS cluster

To create your cluster, you will need the following data:

- name of your EKS cluster,
- AWS region in which you wish to deploy your cluster,
- the amount of nodes you would like to have,
- the amount of [on-demand](#) and [spot](#) instances to use.

Note

[spot](#) instances are not recommended for production environment, but may be useful e.g. for testing purposes.

The most easy and visually clear way is to describe the desired cluster in YAML and to pass this configuration to the `eksctl` command.

The following example configures a EKS cluster with one [managed node group](#):

```
apiVersion: eksctl.io/v1alpha5
kind: ClusterConfig

metadata:
  name: test-cluster
  region: eu-west-2

nodeGroups:
- name: ng-1
  minSize: 3
  maxSize: 5
  instancesDistribution:
    maxPrice: 0.15
    instanceTypes: ["m5.xlarge", "m5.2xlarge"] # At least two instance types should be specified
    onDemandBaseCapacity: 0
    onDemandPercentageAboveBaseCapacity: 50
    spotInstancePools: 2
  tags:
    'iit-billing-tag': 'cloud'
  preBootstrapCommands:
    - "echo 'OPTIONS="--default-ulimit nofile=1048576:1048576"' >> /etc/sysconfig/docker"
    - "systemctl restart docker"
```

Note

`preBootstrapCommands` section is used in the above example to increase the limits for the amount of opened files: this is important and shouldn't be omitted, taking into account the default EKS soft limit of 65536 files.

When the cluster configuration file is ready, you can actually create your cluster by the following command:

```
$ eksctl create cluster -f ~/cluster.yaml
```

11.2.3 Install the Operator

1. Create a namespace and set the context for the namespace. The resource names must be unique within the namespace and provide a way to divide cluster resources between users spread across multiple projects.

So, create the namespace and save it in the namespace context for subsequent commands as follows (replace the `<namespace name>` placeholder with some descriptive name):

```
$ kubectl create namespace <namespace name>
$ kubectl config set-context $(kubectl config current-context) --namespace=<namespace name>
```

At success, you will see the message that namespace/ was created, and the context was modified.

2. Use the following `git clone` command to download the correct branch of the `percona-xtradb-cluster-operator` repository:

```
$ git clone -b v1.12.0 https://github.com/percona/percona-xtradb-cluster-operator
```

After the repository is downloaded, change the directory to run the rest of the commands in this document:

```
$ cd percona-xtradb-cluster-operator
```

3. Deploy the Operator with the following command:

```
$ kubectl apply -f deploy/bundle.yaml
```

The following confirmation is returned:

```
customresourcedefinition.apiextensions.k8s.io/perconaxtradbclusters.pxc.percona.com created
customresourcedefinition.apiextensions.k8s.io/perconaxtradbclusterbackups.pxc.percona.com created
customresourcedefinition.apiextensions.k8s.io/perconaxtradbclusterrestores.pxc.percona.com created
customresourcedefinition.apiextensions.k8s.io/perconaxtradbbackups.pxc.percona.com created
role.rbac.authorization.k8s.io/percona-xtradb-cluster-operator created
serviceaccount/percona-xtradb-cluster-operator created
rolebinding.rbac.authorization.k8s.io/service-account-percona-xtradb-cluster-operator created
deployment.apps/percona-xtradb-cluster-operator created
```

4. The operator has been started, and you can create the Percona XtraDB cluster:

```
$ kubectl apply -f deploy/cr.yaml
```

The process could take some time. The return statement confirms the creation:

```
perconaxtradbcluster.pxc.percona.com/cluster1 created
```

5. During previous steps, the Operator has generated several [secrets](#), including the password for the `root` user, which you will need to access the cluster.

Use `kubectl get secrets` command to see the list of Secrets objects (by default Secrets object you are interested in has `cluster1-secrets` name). Then `kubectl get secret cluster1-secrets -o yaml` will return the YAML file with generated secrets, including the root password which should look as follows:

```
...
data:
  ...
  root: cm9vdF9wYXNzd29yZA==
```

Here the actual password is base64-encoded, and `echo 'cm9vdF9wYXNzd29yZA==' | base64 --decode` will bring it back to a human-readable form (in this example it will be a `root_password` string).

- Now you can check whether you are able to connect to MySQL from the outside with the help of the `kubectl port-forward` command as follows:

```
$ kubectl port-forward svc/example-proxysql 3306:3306 &  
$ mysql -h 127.0.0.1 -P 3306 -uroot -proot_password
```

Last update: 2022-12-04

11.3 Install Percona XtraDB Cluster on Azure Kubernetes Service (AKS)

This guide shows you how to deploy Percona Operator for MySQL based on Percona XtraDB Cluster on Microsoft Azure Kubernetes Service (AKS). The document assumes some experience with the platform. For more information on the AKS, see the [Microsoft AKS official documentation](#).

11.3.1 Prerequisites

The following tools are used in this guide and therefore should be preinstalled:

1. **Azure Command Line Interface (Azure CLI)** for interacting with the different parts of AKS. You can install it following the [official installation instructions for your system](#).
2. **kubectl** to manage and deploy applications on Kubernetes. Install it [following the official installation instructions](#).

Also, you need to sign in with Azure CLI using your credentials according to the [official guide](#).

11.3.2 Create and configure the AKS cluster

To create your cluster, you will need the following data:

- name of your AKS cluster,
- an [Azure resource group](#), in which resources of your cluster will be deployed and managed.
- the amount of nodes you would like to have.

You can create your cluster via command line using `az aks create` command. The following command will create a 3-node cluster named `cluster1` within some [already existing](#) resource group named `my-resource-group`:

```
$ az aks create --resource-group my-resource-group --name cluster1 --enable-managed-identity --node-count 3 --node-vm-size Standard_B4ms --node-osdisk-size 30 --network-plugin kubenet --generate-ssh-keys --outbound-type loadbalancer
```

Other parameters in the above example specify that we are creating a cluster with machine type of [Standard_B4ms](#) and OS disk size reduced to 30 GiB. You can see detailed information about cluster creation options in the [AKS official documentation](#).

You may wait a few minutes for the cluster to be generated.

Now you should configure the command-line access to your newly created cluster to make `kubectl` be able to use it.

```
az aks get-credentials --resource-group my-resource-group --name cluster1
```

11.4 Install the Operator and deploy your Percona XtraDB Cluster

1. Deploy the Operator. By default deployment will be done in the `default` namespace. If that's not the desired one, you can create a new namespace and/or set the context for the namespace as follows (replace the `<namespace name>` placeholder with some descriptive name):

```
$ kubectl create namespace <namespace name>
$ kubectl config set-context $(kubectl config current-context) --namespace=<namespace name>
```

At success, you will see the message that `namespace/<namespace name>` was created, and the context (`<cluster name>`) was modified.

Deploy the Operator using the following command:

```
$ kubectl apply -f https://raw.githubusercontent.com/percona/percona-xtradb-cluster-operator/v1.12.0/deploy/bundle.yaml
```

Expected output

```

customresourcedefinition.apiextensions.k8s.io/perconaxtradbclusters.pxc.percona.com created
customresourcedefinition.apiextensions.k8s.io/perconaxtradbclusterbackups.pxc.percona.com created
customresourcedefinition.apiextensions.k8s.io/perconaxtradbclusterrestores.pxc.percona.com created
customresourcedefinition.apiextensions.k8s.io/perconaxtradbbackups.pxc.percona.com created
role.rbac.authorization.k8s.io/percona-xtradb-cluster-operator created
serviceaccount/percona-xtradb-cluster-operator created
rolebinding.rbac.authorization.k8s.io/service-account-percona-xtradb-cluster-operator created
deployment.apps/percona-xtradb-cluster-operator created

```

2. The operator has been started, and you can deploy Percona XtraDB Cluster:

```
$ kubectl apply -f https://raw.githubusercontent.com/percona/percona-xtradb-cluster-operator/v1.12.0/deploy/cr.yaml
```

Expected output

```
perconaxtradbcluster.pxc.percona.com/ cluster1 created
```

Note

This deploys default Percona XtraDB Cluster configuration with three HAProxy and three XtraDB Cluster instances. Please see [deploy/cr.yaml](#) and [Custom Resource Options](#) for the configuration options. You can clone the repository with all manifests and source code by executing the following command:

```
$ git clone -b v1.12.0 https://github.com/percona/percona-xtradb-cluster-operator
```

After editing the needed options, apply your modified `deploy/cr.yaml` file as follows:

```
$ kubectl apply -f deploy/cr.yaml
```

The creation process may take some time. When the process is over your cluster will obtain the `ready` status. You can check it with the following command:

```
$ kubectl get pxc
```

Expected output

NAME	ENDPOINT	STATUS	PXC	PROXYSQL	HAPROXY	AGE
cluster1	cluster1-haproxy.default	ready	3		3	5m51s

11.4.1 Verifying the cluster operation

It may take ten minutes to get the cluster started. When `kubectl get pxc` command finally shows you the cluster status as `ready`, you can try to connect to the cluster.

1. You will need the login and password for the admin user to access the cluster. Use `kubectl get secrets` command to see the list of Secrets objects (by default the Secrets object you are interested in has `cluster1-secrets` name). You can use the following command to get the password of the `root` user:

```
$ kubectl get secrets cluster1-secrets -o yaml -o jsonpath='{.data.root}' | base64 --decode | tr '\n' ' ' && echo " "
```

2. Run a container with `mysql` tool and connect its console output to your terminal. The following command will do this, naming the new Pod `percona-client`:

```
$ kubectl run -i --rm --tty percona-client --image=percona:8.0 --restart=Never -- bash -il
```

Executing it may require some time to deploy the correspondent Pod.

3. Now run `mysql` tool in the `percona-client` command shell using the password obtained from the secret. The command will look different depending on whether your cluster provides load balancing with [HAProxy](#) (the default choice) or [ProxySQL](#):

with HAProxy (default)

```
$ mysql -h cluster1-haproxy -uroot -proot_password
```

with ProxySQL

```
$ mysql -h cluster1-proxysql -uroot -proot_password
```

11.4.2 Troubleshooting

If `kubectl get pxc` command doesn't show `ready` status too long, you can check the creation process with the `kubectl get pods` command:

```
$ kubectl get pods
```

Expected output

NAME	READY	STATUS	RESTARTS	AGE
cluster1-haproxy-0	2/2	Running	0	6m17s
cluster1-haproxy-1	2/2	Running	0	4m59s
cluster1-haproxy-2	2/2	Running	0	4m36s
cluster1-pxc-0	3/3	Running	0	6m17s
cluster1-pxc-1	3/3	Running	0	5m3s
cluster1-pxc-2	3/3	Running	0	3m56s
percona-xtradb-cluster-operator-79966668bd-rswbk	1/1	Running	0	9m54s

If the command output had shown some errors, you can examine the problematic Pod with the `kubectl describe <pod name>` command as follows:

```
$ kubectl describe pod cluster1-pxc-2
```

Review the detailed information for `Warning` statements and then correct the configuration. An example of a warning is as follows:

```
Warning FailedScheduling 68s (x4 over 2m22s) default-scheduler 0/1 nodes are available: 1 node(s) didn't match pod affinity/anti-affinity, 1 node(s) didn't satisfy existing pods anti-affinity rules.
```

11.4.3 Removing the AKS cluster

To delete your cluster, you will need the following data:

- name of your AKS cluster,
- AWS region in which you have deployed your cluster.

You can clean up the cluster with the `az aks delete` command as follows (with real names instead of `<resource group>` and `<cluster name>` placeholders):

```
$ az aks delete --name <cluster name> --resource-group <resource group> --yes --no-wait
```

It may take ten minutes to get the cluster actually deleted after executing this command.

 **Warning**

After deleting the cluster, all data stored in it will be lost!

Last update: 2022-12-06

11.5 Install Percona XtraDB Cluster on Kubernetes

1. First of all, clone the percona-xtradb-cluster-operator repository:

```
$ git clone -b v1.12.0 https://github.com/percona/percona-xtradb-cluster-operator
$ cd percona-xtradb-cluster-operator
```

 **Note**

It is crucial to specify the right branch with `-b` option while cloning the code on this step. Please be careful.

1. Now Custom Resource Definition for Percona XtraDB Cluster should be created from the `deploy/crd.yaml` file. Custom Resource Definition extends the standard set of resources which Kubernetes “knows” about with the new items (in our case ones which are the core of the operator).

This step should be done only once; it does not need to be repeated with the next Operator deployments, etc.

```
$ kubectl apply -f deploy/crd.yaml
```

2. The next thing to do is to add the `pxc` namespace to Kubernetes, not forgetting to set the correspondent context for further steps:

```
$ kubectl create namespace pxc
$ kubectl config set-context $(kubectl config current-context) --namespace=pxc
```

3. Now RBAC (role-based access control) for Percona XtraDB Cluster should be set up from the `deploy/rbac.yaml` file. Briefly speaking, role-based access is based on specifically defined roles and actions corresponding to them, allowed to be done on specific Kubernetes resources (details about users and roles can be found in [Kubernetes documentation](#)).

```
$ kubectl apply -f deploy/rbac.yaml
```

Note

Setting RBAC requires your user to have cluster-admin role privileges. For example, those using Google Kubernetes Engine can grant user needed privileges with the following command:

```
$ kubectl create clusterrolebinding cluster-admin-binding --clusterrole=cluster-admin --user=$(gcloud config get-value core/account)
```

Finally it's time to start the operator within Kubernetes:

```
$ kubectl apply -f deploy/operator.yaml
```

Note

You can simplify the Operator installation by applying a single `deploy/bundle.yaml` file instead of running commands from the steps 2 and 4:

```
$ kubectl apply -f deploy/bundle.yaml
```

This will automatically create Custom Resource Definition, set up role-based access control and install the Operator as one single action.

4. Now that's time to add the Percona XtraDB Cluster users [Secrets](#) with logins and passwords to Kubernetes. By default, the Operator generates users Secrets automatically, and *no actions are required at this step*.

Still, you can generate and apply your Secrets by your own. In this case, place logins and plaintext passwords for the user accounts in the data section of the `deploy/secrets.yaml` file; after editing is finished, create users Secrets with the following command:

```
$ kubectl create -f deploy/secrets.yaml
```

More details about secrets can be found in [Users](#).

5. Now certificates should be generated. By default, the Operator generates certificates automatically, and *no actions are required at this step*. Still, you can generate and apply your own certificates as secrets according to the [TLS instructions](#).
6. After the operator is started and user secrets are added, Percona XtraDB Cluster can be created at any time with the following command:

```
$ kubectl apply -f deploy/cr.yaml
```

Creation process will take some time. The process is over when both operator and replica set pod have reached their Running status:

NAME	READY	STATUS	RESTARTS	AGE
cluster1-haproxy-0	2/2	Running	0	6m17s
cluster1-haproxy-1	2/2	Running	0	4m59s
cluster1-haproxy-2	2/2	Running	0	4m36s
cluster1-pxc-0	3/3	Running	0	6m17s
cluster1-pxc-1	3/3	Running	0	5m3s
cluster1-pxc-2	3/3	Running	0	3m56s
percona-xtradb-cluster-operator-7996668bd-rswbk	1/1	Running	0	9m54s

7. Check connectivity to newly created cluster

```
$ kubectl run -i --rm --tty percona-client --image=percona:8.0 --restart=Never -- bash -il
percona-client:/$ mysql -h cluster1-haproxy -uroot -proot_password
```

This command will connect you to the MySQL monitor.

```
mysql: [Warning] Using a password on the command line interface can be insecure.
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1976
Server version: 8.0.19-10 Percona XtraDB Cluster (GPL), Release rel10, Revision 727f180, WSREP version 26.4.3

Copyright (c) 2009-2020 Percona LLC and/or its affiliates
Copyright (c) 2000, 2020, Oracle and/or its affiliates. All rights reserved.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
```

Last update: 2022-12-07

11.6 Install Percona XtraDB Cluster on OpenShift

Percona Operator for Percona XtraDB Cluster is a [Red Hat Certified Operator](#). This means that Percona Operator is portable across hybrid clouds and fully supports the Red Hat OpenShift lifecycle.

Installing Percona XtraDB Cluster on OpenShift includes two steps:

- Installing the Percona Operator for MySQL,
- Install Percona XtraDB Cluster using the Operator.

11.6.1 Install the Operator

You can install Percona Operator for MySQL on OpenShift using the [Red Hat Marketplace](#) web interface or using the command line interface.

Install the Operator via the Red Hat Marketplace

1. login to the Red Hat Marketplace and register your cluster [following the official instructions](#).
2. Go to the [Percona Operator for MySQL](#) page and click the Free trial button:

Percona Kubernetes Operator for Percona XtraDB Cluster

By Percona

The Operator is an open-source drop in replacement for MySQL Enterprise with synchronous replication running on Kubernetes. It automates the deployment and management of the members in your Percona XtraDB Cluster environment.

Software version	Runs on	Delivery method	Rating
1.6.0	OpenShift 4.3	Operator	☆☆☆☆☆ Not rated

[Overview](#)

[Documentation](#)

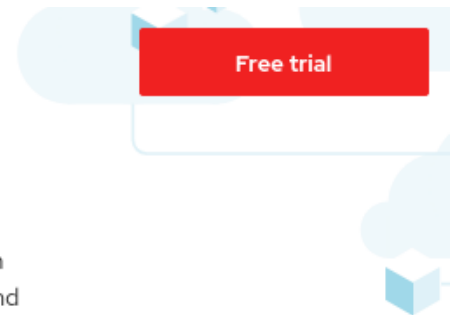
[Pricing](#)

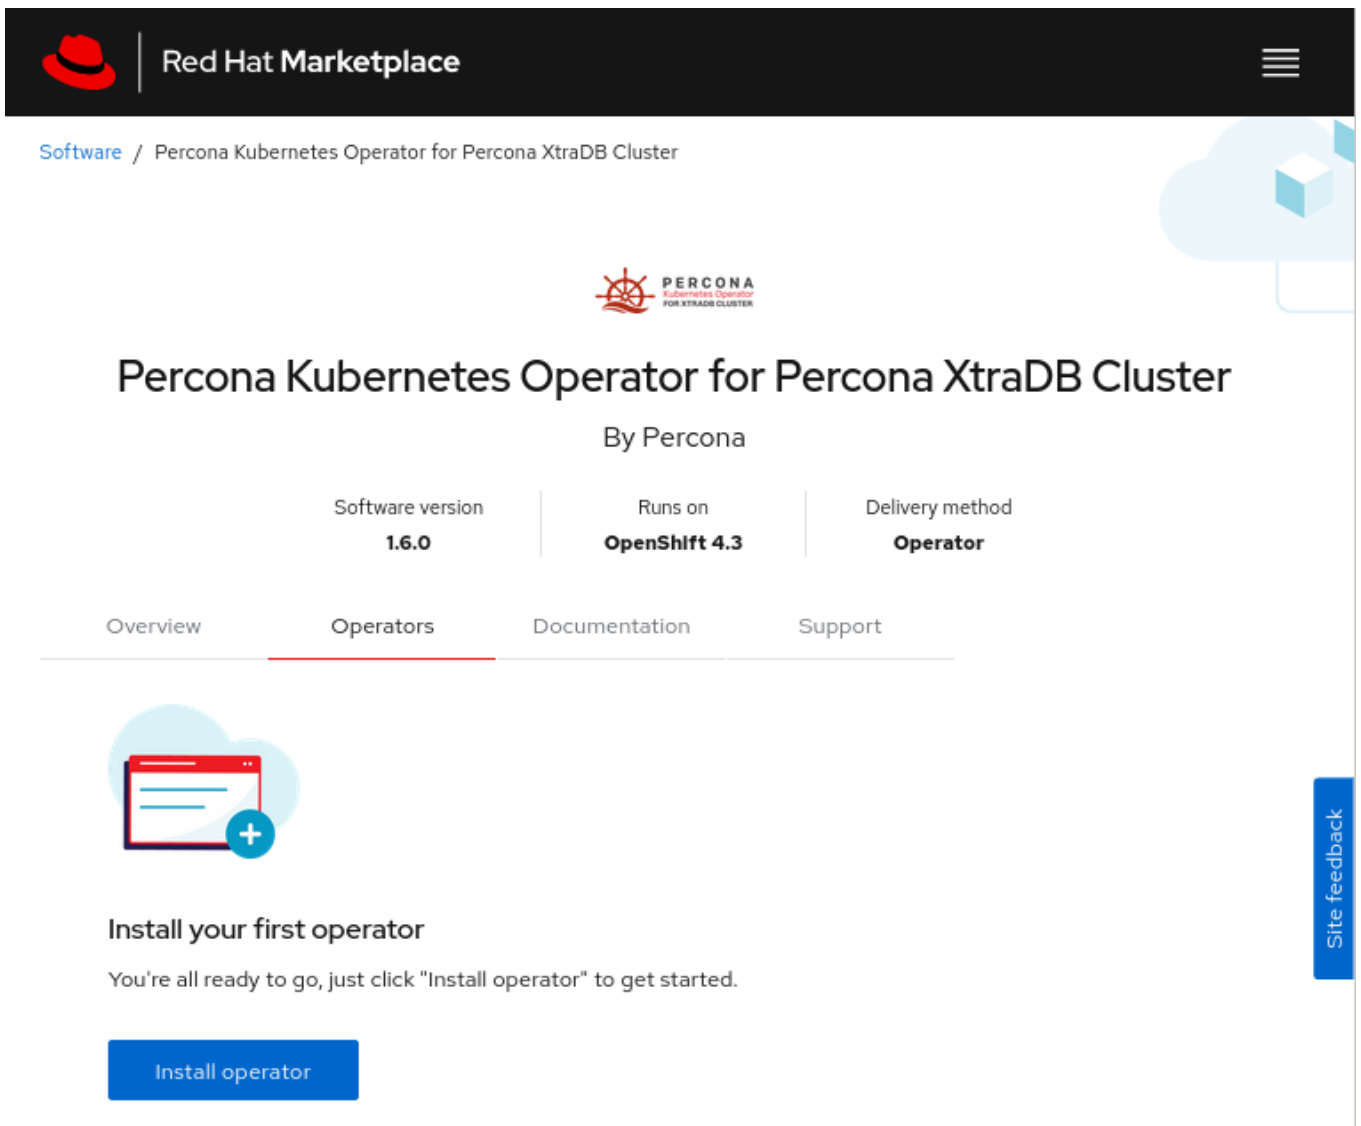
[Help](#)

Based on our best practices for deployment and configuration the Operator contains everything you need to quickly and consistently deploy and scale XtraDB Cluster into a Kubernetes cluster. The Operator is Red Hat OpenShift Certified and is available in concurrent release with our software products.

Here you can “start trial” of the Operator for 0.0 USD.

3. When finished, chose [Workspace->Software](#) in the system menu on the top and choose the Operator:





Red Hat Marketplace

Software / Percona Kubernetes Operator for Percona XtraDB Cluster

PERCONA
Kubernetes Operator
FOR XTRADB CLUSTER

Percona Kubernetes Operator for Percona XtraDB Cluster

By Percona

Software version 1.6.0	Runs on OpenShift 4.3	Delivery method Operator
----------------------------------	---------------------------------	------------------------------------

Overview Operators Documentation Support

Install your first operator

You're all ready to go, just click "Install operator" to get started.

[Install operator](#)

Site feedback

Click the `Install Operator` button.

Install the Operator via the command-line interface

1. Clone the `percona-xtradb-cluster-operator` repository:

```
$ git clone -b v1.12.0 https://github.com/percona/percona-xtradb-cluster-operator
$ cd percona-xtradb-cluster-operator
```

Note

It is crucial to specify the right branch with the `-b` option while cloning the code on this step. Please be careful.

1. Now Custom Resource Definition for Percona XtraDB Cluster should be created from the `deploy/crd.yaml` file. Custom Resource Definition extends the standard set of resources which Kubernetes “knows” about with the new items (in our case ones which are the core of the operator).

This step should be done only once; it does not need to be repeated with the next Operator deployments, etc.

```
$ oc apply -f deploy/crd.yaml
```

Note

Setting Custom Resource Definition requires your user to have cluster-admin role privileges.

If you want to manage your Percona XtraDB Cluster with a non-privileged user, necessary permissions can be granted by applying the next clusterrole:

```
$ oc create clusterrole pxc-admin --verb="" --resource=perconaxtradbclusters.pxc.percona.com,perconaxtradbclusters.pxc.percona.com/status,perconaxtradbclusterbackups.pxc.percona.com,perconaxtradbclusterbackups.pxc.percona.com/status,perconaxtradbclusterrestores.pxc.percona.com,perconaxtradbclusterrestores.pxc.percona.com/status
$ oc adm policy add-cluster-role-to-user pxc-admin <some-user>
```

If you have a [cert-manager](#) installed, then you have to execute two more commands to be able to manage certificates with a non-privileged user:

```
$ oc create clusterrole cert-admin --verb="" --resource=issuers.certmanager.k8s.io,certificates.certmanager.k8s.io
$ oc adm policy add-cluster-role-to-user cert-admin <some-user>
```

2. The next thing to do is to create a new `pxc` project:

```
$ oc new-project pxc
```

3. Now RBAC (role-based access control) for Percona XtraDB Cluster should be set up from the `deploy/rbac.yaml` file. Briefly speaking, role-based access is based on specifically defined roles and actions corresponding to them, allowed to be done on specific Kubernetes resources (details about users and roles can be found in [OpenShift documentation](#)).

```
$ oc apply -f deploy/rbac.yaml
```

Finally, it's time to start the operator within OpenShift:

```
$ oc apply -f deploy/operator.yaml
```

Note

You can simplify the Operator installation by applying a single `deploy/bundle.yaml` file instead of running commands from the steps 2 and 4:

```
$ oc apply -f deploy/bundle.yaml
```

This will automatically create Custom Resource Definition, set up role-based access control and install the Operator as one single action.

11.6.2 Install Percona XtraDB Cluster

1. Now that's time to add the Percona XtraDB Cluster users [Secrets](#) with logins and passwords to Kubernetes. By default, the Operator generates users Secrets automatically, and *no actions are required at this step*.

Still, you can generate and apply your Secrets by your own. In this case, place logins and plaintext passwords for the user accounts in the data section of the `deploy/secrets.yaml` file; after editing is finished, create users Secrets with the following command:

```
$ oc create -f deploy/secrets.yaml
```

More details about secrets can be found in [Users](#).

- Now certificates should be generated. By default, the Operator generates certificates automatically, and no actions are required at this step. Still, you can generate and apply your own certificates as secrets according to the [TLS instructions](#).
- After the operator is started and user secrets are added, Percona XtraDB Cluster can be created at any time with the following command:

```
$ oc apply -f deploy/cr.yaml
```

Creation process will take some time. The process is over when both operator and replica set pod have reached their Running status:

NAME	READY	STATUS	RESTARTS	AGE
cluster1-haproxy-0	2/2	Running	0	6m17s
cluster1-haproxy-1	2/2	Running	0	4m59s
cluster1-haproxy-2	2/2	Running	0	4m36s
cluster1-pxc-0	3/3	Running	0	6m17s
cluster1-pxc-1	3/3	Running	0	5m3s
cluster1-pxc-2	3/3	Running	0	3m56s
percona-xtradb-cluster-operator-7996668bd-rswbk	1/1	Running	0	9m54s

- Check connectivity to newly created cluster. Run a container with MySQL monitor and connect its console output to your terminal. The following command will do this, naming the new Pod `percona-client`:

```
$ oc run -i --rm --tty percona-client --image=percona:8.0 --restart=Never -- bash -il
percona-client:/$ mysql -h cluster1-haproxy -uroot -proot_password
```

This command will connect you to the MySQL monitor.

```
mysql: [Warning] Using a password on the command line interface can be insecure.
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1976
Server version: 8.0.19-10 Percona XtraDB Cluster (GPL), Release rel10, Revision 727f180, WSREP version 26.4.3

Copyright (c) 2009-2020 Percona LLC and/or its affiliates
Copyright (c) 2000, 2020, Oracle and/or its affiliates. All rights reserved.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
```

Last update: 2022-12-07

11.7 Use docker images from a custom registry

Using images from a private Docker registry may be useful in different situations: it may be related to storing images inside of a company, for privacy and security reasons, etc. In such cases, Percona Distribution for MySQL Operator based on Percona

XtraDB Cluster allows to use a custom registry, and the following instruction illustrates how this can be done by the example of the Operator deployed in the OpenShift environment.

1. First of all login to the OpenShift and create project.

```
$ oc login
Authentication required for https://192.168.1.100:8443 (openshift)
Username: admin
Password:
Login successful.
$ oc new-project pxc
Now using project "pxc" on server "https://192.168.1.100:8443".
```

2. There are two things you will need to configure your custom registry access:

- the token for your user
- your registry IP address.

The token can be find out with the following command:

```
$ oc whoami -t
AD08CqCDappWR4hxjfdqwijEHei31yXAvWg61Jg210s
```

And the following one tells you the registry IP address:

```
$ kubectl get services/docker-registry -n default
NAME          TYPE        CLUSTER-IP   EXTERNAL-IP   PORT(S)    AGE
docker-registry ClusterIP   172.30.162.173 <none>        5000/TCP   1d
```

3. Now you can use the obtained token and address to login to the registry:

```
$ docker login -u admin -p AD08CqCDappWR4hxjfdqwijEHei31yXAvWg61Jg210s 172.30.162.173:5000
Login Succeeded
```

4. Pull the needed image by its SHA digest:

```
$ docker pull docker.io/perconalab/percona-xtradb-cluster-operator@sha256:841c07eef30605080bfe80e549f9332ab6b9755fcbc42aacbf86e4ac9ef0e444
Trying to pull repository docker.io/perconalab/percona-xtradb-cluster-operator ...
sha256:841c07eef30605080bfe80e549f9332ab6b9755fcbc42aacbf86e4ac9ef0e444: Pulling from docker.io/perconalab/percona-xtradb-cluster-operator
Digest: sha256:841c07eef30605080bfe80e549f9332ab6b9755fcbc42aacbf86e4ac9ef0e444
Status: Image is up to date for docker.io/perconalab/percona-xtradb-cluster-operator@sha256:841c07eef30605080bfe80e549f9332ab6b9755fcbc42aacbf86e4ac9ef0e444
```

You can find correct names and SHA digests in the [current list of the Operator-related images officially certified by Percona](#).

5. The following way is used to push an image to the custom registry (into the OpenShift pxc project):

```
$ docker tag \
  docker.io/perconalab/percona-xtradb-cluster-operator@sha256:841c07eef30605080bfe80e549f9332ab6b9755fcbc42aacbf86e4ac9ef0e444 \
  172.30.162.173:5000/pxc/percona-xtradb-cluster-operator:1.12.0
$ docker push 172.30.162.173:5000/pxc/percona-xtradb-cluster-operator:1.12.0
```

6. Check the image in the OpenShift registry with the following command:

```
$ oc get is
NAME          DOCKER REPO          TAGS    UPDATED
percona-xtradb-cluster-operator  docker-registry.default.svc:5000/pxc/percona-xtradb-cluster-operator  1.12.0  2 hours ago
```

7. When the custom registry image is Ok, put a Docker Repo + Tag string (it should look like `docker-registry.default.svc:5000/pxc/percona-xtradb-cluster-operator:1.12.0`) into the `initImage` option in `deploy/operator.yaml` configuration file.
8. Repeat steps 3-5 for other images, updating the `image\` options in the corresponding sections of the the ``deploy/cr.yaml`` file.`

Note

Don't forget to set `upgradeoptions.apply` option to `Disabled`. Otherwise `Smart Upgrade functionality` will try using the image recommended by the Version Service instead of the custom one.

Please note it is possible to specify `imagePullSecrets` option for the images, if the registry requires authentication.

9. Now follow the standard [Percona Operator for MySQL installation instruction](#).

Last update: 2022-12-04

12. Configuration

12.1 Users

MySQL user accounts within the Cluster can be divided into two different groups:

- *application-level users*: the unprivileged user accounts,
- *system-level users*: the accounts needed to automate the cluster deployment and management tasks, such as Percona XtraDB Cluster Health checks or ProxySQL integration.

As these two groups of user accounts serve different purposes, they are considered separately in the following sections.

12.1.1 Unprivileged users

There are no unprivileged (general purpose) user accounts created by default. If you need general purpose users, please run commands below:

```
$ kubectl run -it --rm percona-client --image=percona:8.0 --restart=Never -- mysql -hcluster1-pxc -uroot -proot_password
mysql> GRANT ALL PRIVILEGES ON database1.* TO 'user1'@'%' IDENTIFIED BY 'password1';
```

Note

MySQL password here should not exceed 32 characters due to the [replication-specific limit introduced in MySQL 5.7.5](#).

Verify that the user was created successfully. If successful, the following command will let you successfully login to MySQL shell via ProxySQL:

```
$ kubectl run -it --rm percona-client --image=percona:8.0 --restart=Never -- bash -il
percona-client:/$ mysql -h cluster1-proxysql -uuser1 -ppassword1
mysql> SELECT * FROM database1.table1 LIMIT 1;
```

You may also try executing any simple SQL statement to ensure the permissions have been successfully granted.

12.1.2 System Users

To automate the deployment and management of the cluster components, the Operator requires system-level Percona XtraDB Cluster users.

Credentials for these users are stored as a [Kubernetes Secrets](#) object. The Operator requires Kubernetes Secrets before Percona XtraDB Cluster is started. It will either use existing Secrets or create a new Secrets object with randomly generated passwords if it didn't exist. The name of the required Secrets (`cluster1` by default) should be set in the `spec.secretsName` option of the `deploy/cr.yaml` configuration file.

The following table shows system users' names and purposes.

Warning

These users should not be used to run an application.

User Purpose	Username	Password Secret Key	Description
Admin	root	root	Database administrative user, can be used by the application if needed
ProxySQLAdmin	proxyadmin	proxyadmin	ProxySQL administrative user, can be used to add general-purpose ProxySQL users
Backup	xtrabackup	xtrabackup	User to run backups
Cluster Check	clustercheck	clustercheck	User for liveness checks and readiness checks
Monitoring	monitor	monitor	User for internal monitoring purposes and PMM agent
PMM Server Password	should be set through the operator options	pmmserver	Password used to access PMM Server. Password-based authorization method is deprecated since the Operator 1.11.0. Use token-based authorization instead
Operator Admin	operator	operator	Database administrative user, should be used only by the Operator
Replication	replication	replication	Administrative user needed for cross-site Percona XtraDB Cluster

YAML Object Format

The default name of the Secrets object for these users is `cluster1-secrets` and can be set in the CR for your cluster in `spec.secretName` to something different. When you create the object yourself, it should match the following simple format:

```
apiVersion: v1
kind: Secret
metadata:
  name: cluster1-secrets
type: Opaque
stringData:
  root: root_password
  xtrabackup: backup_password
  monitor: monitory
  clustercheck: clustercheckpassword
  proxyadmin: admin_password
  pmmserver: admin
  operator: operatoradmin
  replication: repl_password
```

The example above matches what is shipped in `deploy/secrets.yaml` which contains default passwords. You should NOT use these in production, but they are present to assist in automated testing or simple use in a development environment.

As you can see, because we use the `stringData` type when creating the Secrets object, all values for each key/value pair are stated in plain text format convenient from the user's point of view. But the resulting Secrets object contains passwords stored as `data` - i.e., base64-encoded strings. If you want to update any field, you'll need to encode the value into base64 format. To do this, you can run `echo -n "password" | base64 --wrap=0` (or just `echo -n "password" | base64` in case of Apple macOS) in your local shell to get valid values. For example, setting the PMM Server user's password to `new_password` in the `cluster1-secrets` object can be done with the following command:

in Linux

```
$ kubectl patch secret/cluster1-secrets -p '{"data":{"pmmserver": "'$(echo -n new_password | base64 --wrap=0)'"}}'
```

in macOS

```
$ kubectl patch secret/cluster1-secrets -p '{"data":{"pmmserver": "'$(echo -n new_password | base64)'"}}'
```

Password Rotation Policies and Timing

When there is a change in user secrets, the Operator creates the necessary transaction to change passwords. This rotation happens almost instantly (the delay can be up to a few seconds), and it's not needed to take any action beyond changing the password.

Note

Please don't change `secretName` option in CR, make changes inside the secrets object itself.

Marking System Users In MySQL

Starting with MySQL 8.0.16, a new feature called Account Categories has been implemented, which allows us to mark our system users as such. See [the official documentation on this feature](#) for more details.

12.1.3 Development Mode

To make development and testing easier, `deploy/secrets.yaml` secrets file contains default passwords for Percona XtraDB Cluster system users.

These development mode credentials from `deploy/secrets.yaml` are:

Secret Key	Secret Value
root	root_password
xtrabackup	backup_password
monitor	monitor
clustercheck	clustercheckpassword
proxyuser	s3cret
proxyadmin	admin_password
pmmserver	admin
operator	operatoradmin
replication	repl_password

Warning

Do not use the default Percona XtraDB Cluster user passwords in production!

Last update: 2022-12-07

12.2 Changing MySQL Options

You may require a configuration change for your application. MySQL allows the option to configure the database with a configuration file. You can pass options from the `my.cnf` configuration file to be included in the MySQL configuration in one of the following ways:

- edit the `deploy/cr.yaml` file,
- use a ConfigMap,
- use a Secret object.

Often there's no need to add custom options, as the Operator takes care of providing MySQL with reasonable defaults. Also, some MySQL options can not be changed: you shouldn't change `require_secure_transport` option to `ON`, as it would break the behavior of the Operator.

Note

If you still need something equal to `require_secure_transport=ON` to force encrypted connections between client and server, the most convenient workaround would be [creating MySQL users with REQUIRE SSL option](#).

12.2.1 Edit the `deploy/cr.yaml` file

You can add options from the `my.cnf` configuration file by editing the configuration section of the `deploy/cr.yaml`. Here is an example:

```
spec:
  secretsName: cluster1-secrets
  pxc:
    ...
    configuration: |
      [mysqld]
      wsrep_debug=CLIENT
      [sst]
      wsrep_debug=CLIENT
```

See the [Custom Resource options, PXC section](#) for more details.

12.2.2 Use a ConfigMap

You can use a configmap and the cluster restart to reset configuration options. A configmap allows Kubernetes to pass or update configuration data inside a containerized application.

Use the `kubectl` command to create the configmap from external resources, for more information see [Configure a Pod to use a ConfigMap](#).

For example, let's suppose that your application requires more connections. To increase your `max_connections` setting in MySQL, you define a `my.cnf` configuration file with the following setting:

```
[mysqld]
...
max_connections=250
```

You can create a configmap from the `my.cnf` file with the `kubectl create configmap` command.

You should use the combination of the cluster name with the `-pxc` suffix as the naming convention for the configmap. To find the cluster name, you can use the following command:

```
$ kubectl get pxc
```

The syntax for `kubectl create configmap` command is:

```
$ kubectl create configmap <cluster-name>-pxc <resource-type>=<resource-name>
```

The following example defines `cluster1-pxc` as the configmap name and the `my.cnf` file as the data source:

```
$ kubectl create configmap cluster1-pxc --from-file=my.cnf
```

To view the created configmap, use the following command:

```
$ kubectl describe configmaps cluster1-pxc
```

12.2.3 Use a Secret Object

The Operator can also store configuration options in [Kubernetes Secrets](#). This can be useful if you need additional protection for some sensitive data.

You should create a Secret object with a specific name, composed of your cluster name and the `pxc` suffix.

Note

To find the cluster name, you can use the following command:

```
$ kubectl get pxc
```

Configuration options should be put inside a specific key inside of the `data` section. The name of this key is `my.cnf` for Percona XtraDB Cluster Pods.

Actual options should be encoded with [Base64](#).

For example, let's define a `my.cnf` configuration file and put there a pair of MySQL options we used in the previous example:

```
[mysqld]
wsrep_debug=CLIENT
[ssst]
wsrep_debug=CLIENT
```

You can get a Base64 encoded string from your options via the command line as follows:

in Linux

```
$ cat my.cnf | base64 --wrap=0
```

in macOS

```
$ cat my.cnf | base64
```

Note

Similarly, you can read the list of options from a Base64 encoded string:

```
$ echo "W215c3FsZF0kd3NyZXBfZGVidWc9T04KW3NzdF0kd3NyZXBfZGVidWc9T04K" | base64 --decode
```

Finally, use a yml file to create the Secret object. For example, you can create a `deploy/my-pxc-secret.yml` file with the following contents:

```
apiVersion: v1
kind: Secret
metadata:
  name: cluster1-pxc
data:
  my.cnf: "W215c3FsZF0kd3NyZXBfZGVidWc9T04KW3NzdF0kd3NyZXBfZGVidWc9T04K"
```

When ready, apply it with the following command:

```
$ kubectl create -f deploy/my-pxc-secret.yaml
```

Note

Do not forget to restart Percona XtraDB Cluster to ensure the cluster has updated the configuration.

12.2.4 Make changed options visible to Percona XtraDB Cluster

Do not forget to restart Percona XtraDB Cluster to ensure the cluster has updated the configuration (see details on how to connect in the [Install Percona XtraDB Cluster on Kubernetes](#) page).

12.2.5 Auto-tuning MySQL options

Few configuration options for MySQL can be calculated and set by the Operator automatically based on the available Pod resource limits (memory and CPU) **if constant values for these options are not specified by user** (either in CR.yaml or in ConfigMap).

Options which can be set automatically are the following ones:

- `innodb_buffer_pool_size`
- `max_connections`

If Percona XtraDB Cluster Pod limits are defined, then limits values are used to calculate these options. If Percona XtraDB Cluster Pod limits are not defined, auto-tuning is not done.

Also, starting from the Operator 1.12.0, there is another way of auto-tuning. You can use `""` as a value in `spec.pxc.configuration` as follows:

```
pxc:
  configuration: |
    [mysqld]
    innodb_buffer_pool_size={{containerMemoryLimit * 3 / 4}}
    ...
```

Last update: 2022-12-07

12.3 Binding Percona XtraDB Cluster components to Specific Kubernetes/OpenShift Nodes

The operator does good job automatically assigning new Pods to nodes with sufficient to achieve balanced distribution across the cluster. Still there are situations when it worth to ensure that pods will land on specific nodes: for example, to get speed advantages of the SSD equipped machine, or to reduce costs choosing nodes in a same availability zone.

Appropriate sections of the [deploy/cr.yaml](#) file (such as `pxc`, `haproxy`, and `proxysql`) contain keys which can be used to do this, depending on what is the best for a particular situation.

12.3.1 Node selector

`nodeSelector` contains one or more key-value pairs. If the node is not labeled with each key-value pair from the Pod's `nodeSelector`, the Pod will not be able to land on it.

The following example binds the Pod to any node having a self-explanatory `disktype: ssd` label:

```
nodeSelector:
  disktype: ssd
```

12.3.2 Affinity and anti-affinity

Affinity makes Pod eligible (or not eligible - so called "anti-affinity") to be scheduled on the node which already has Pods with specific labels. Particularly this approach is good to to reduce costs making sure several Pods with intensive data exchange will occupy the same availability zone or even the same node - or, on the contrary, to make them land on different nodes or even different availability zones for the high availability and balancing purposes.

Percona Operator for MySQL provides two approaches for doing this:

- simple way to set anti-affinity for Pods, built-in into the Operator,
- more advanced approach based on using standard Kubernetes constraints.

Simple approach - use `topologyKey` of the Percona Operator for MySQL

Percona Operator for MySQL provides a `topologyKey` option, which may have one of the following values:

- `kubernetes.io/hostname` - Pods will avoid residing within the same host,
- `failure-domain.beta.kubernetes.io/zone` - Pods will avoid residing within the same zone,
- `failure-domain.beta.kubernetes.io/region` - Pods will avoid residing within the same region,
- `none` - no constraints are applied.

The following example forces Percona XtraDB Cluster Pods to avoid occupying the same node:

```
affinity:
  topologyKey: "kubernetes.io/hostname"
```

Advanced approach - use standard Kubernetes constraints

Previous way can be used with no special knowledge of the Kubernetes way of assigning Pods to specific nodes. Still in some cases more complex tuning may be needed. In this case `advanced` option placed in the `deploy/cr.yaml` file turns off the effect of the `topologyKey` and allows to use standard Kubernetes affinity constraints of any complexity:

```
affinity:
  advanced:
    podAffinity:
      requiredDuringSchedulingIgnoredDuringExecution:
      - labelSelector:
          matchExpressions:
          - key: security
            operator: In
            values:
```

```

- S1
  topologyKey: failure-domain.beta.kubernetes.io/zone
podAntiAffinity:
  preferredDuringSchedulingIgnoredDuringExecution:
  - weight: 100
    podAffinityTerm:
      labelSelector:
        matchExpressions:
        - key: security
          operator: In
          values:
          - S2
      topologyKey: kubernetes.io/hostname
nodeAffinity:
  requiredDuringSchedulingIgnoredDuringExecution:
    nodeSelectorTerms:
    - matchExpressions:
      - key: kubernetes.io/e2e-az-name
        operator: In
        values:
        - e2e-az1
        - e2e-az2
  preferredDuringSchedulingIgnoredDuringExecution:
  - weight: 1
    preference:
      matchExpressions:
      - key: another-node-label-key
        operator: In
        values:
        - another-node-label-value

```

See explanation of the advanced affinity options in [Kubernetes documentation](#).

12.3.3 Tolerations

Tolerations allow Pods having them to be able to land onto nodes with matching *taints*. Tolerations are expressed as a `key` with an `operator`, which is either `exists` or `equal` (the latter variant also requires a `value` the key is equal to). Moreover, tolerations should have a specified `effect`, which may be a self-explanatory `NoSchedule`, less strict `PreferNoSchedule`, or `NoExecute`. The last variant means that if a *taint* with `NoExecute` is assigned to a node, then any Pod not tolerating this *taint* will be removed from the node, immediately or after the `tolerationSeconds` interval, like in the following example:

```

tolerations:
- key: "node.alpha.kubernetes.io/unreachable"
  operator: "Exists"
  effect: "NoExecute"
  tolerationSeconds: 6000

```

The [Kubernetes Taints and Tolerations](#) contains more examples on this topic.

12.3.4 Priority Classes

Pods may belong to some *priority classes*. This allows the scheduler to distinguish more and less important Pods to resolve the situation when some higher priority Pod cannot be scheduled without evicting a lower priority one. This can be done by adding one or more `PriorityClasses` in your Kubernetes cluster, and specifying the `PriorityClassName` in the `deploy/cr.yaml` file:

```
priorityClassName: high-priority
```

See the [Kubernetes Pods Priority and Preemption documentation](#) to find out how to define and use priority classes in your cluster.

12.3.5 Pod Disruption Budgets

Creating the *Pod Disruption Budget* is the Kubernetes style to limit the number of Pods of an application that can go down simultaneously due to such *voluntary disruptions* as cluster administrator's actions during the update of deployments or nodes, etc. By such a way *Disruption Budgets* allow large applications to retain their high availability while maintenance and other administrative activities.

We recommend to apply *Pod Disruption Budgets* manually to avoid situations when Kubernetes stopped all your database Pods. See the [official Kubernetes documentation](#) for details.

Last update: 2022-09-22

12.4 Labels and annotations

[Labels](#) and [annotations](#) are used to attach additional metadata information to Kubernetes resources.

Labels and annotations are rather similar. The difference between them is that labels are used by Kubernetes to identify and select objects, while annotations are assigning additional *non-identifying* information to resources. Therefore, typical role of Annotations is facilitating integration with some external tools.

12.4.1 Setting labels and annotations in the Custom Resource

You can set labels and/or annotations as key/value string pairs in the Custom Resource metadata section of the `deploy/cr.yaml` as follows:

```
apiVersion: pxc.percona.com/v1
kind: PerconaKtraDBCluster
metadata:
  name: cluster1
  annotations:
    percona.com/issue-vault-token: "true"
  labels:
    ...
```

The easiest way to check which labels are attached to a specific object with is using the additional `--show-labels` option of the `kubectl get` command. Checking the annotations is not much more difficult: it can be done as in the following example:

```
$ kubectl get pod cluster1-pxc-0 -o jsonpath='{.metadata.annotations}'
```

12.4.2 Specifying labels and annotations ignored by the Operator

Sometimes various Kubernetes flavors can add their own annotations to the objects managed by the Operator.

The Operator keeps track of all changes to its objects and can remove annotations that appeared without its participation.

If there are no annotations or labels in the Custom Resource, the Operator does nothing if new label or annotation added to the object.

If there is an annotation or a label specified in the Custom Resource, the Operator starts to manage annotations and labels. In this case it removes unknown annotations and labels.

Still, it is possible to specify which annotations and labels should be ignored by the Operator by listing them in the `spec.ignoreAnnotations` or `spec.ignoreLabels` keys of the `deploy/cr.yaml`, as follows:

```
spec:
  ignoreAnnotations:
    - some.custom.cloud.annotation/smith
  ignoreLabels:
    - some.custom.cloud.label/smith
  ...
```

The Operator will ignore any annotation and label which keys **starts** with the mentioned above examples. For example, the following annotations and labels will be ignored after applying the above `cr.yaml` fragment:

```
annotations:
  some.custom.cloud.annotation/smith: somethinghere
labels:
  some.custom.cloud.label/smith: somethinghere
```

The Operator will ignore specified annotations and labels for all objects: Pods, Services, etc.

Last update: 2022-12-07

12.5 Local Storage support for the Percona Operator for MySQL

Among the wide range of volume types, available in Kubernetes, there are some which allow Pod containers to access part of the local filesystem on the node. Two such options provided by Kubernetes itself are *emptyDir* and *hostPath* volumes. More comprehensive setups require additional components, such as [OpenEBS Container Attached Storage solution](#)

12.5.1 emptyDir

The name of this option is self-explanatory. When Pod having an [emptyDir volume](#) is assigned to a Node, a directory with the specified name is created on this node and exists until this Pod is removed from the node. When the Pod have been deleted, the directory is deleted too with all its content. All containers in the Pod which have mounted this volume will gain read and write access to the correspondent directory.

The `emptyDir` options in the [deploy/cr.yaml](#) file can be used to turn the emptyDir volume on by setting the directory name.

12.5.2 hostPath

A [hostPath volume](#) mounts some existing file or directory from the node's filesystem into the Pod.

The `volumeSpec.hostPath` subsection in the [deploy/cr.yaml](#) file may include `path` and `type` keys to set the node's filesystem object path and to specify whether it is a file, a directory, or something else (e.g. a socket):

```
volumeSpec:
  hostPath:
    path: /data
    type: Directory
```

Please note, that `hostPath` directory is not created automatically! It should be [created manually on the node's filesystem](#). Also, it should have the attributes (access permissions, ownership, SELinux security context) which would allow Pod to access the correspondent filesystem objects according to [pxc.containerSecurityContext](#) and [pxc.podSecurityContext](#).

`hostPath` is useful when you are able to perform manual actions during the first run and have strong need in improved disk performance. Also, please consider using tolerations to avoid cluster migration to different hardware in case of a reboot or a hardware failure.

More details can be found in the [official hostPath Kubernetes documentation](#).

12.5.3 OpenEBS Local Persistent Volume Hostpath

Both *emptyDir* and *hostPath* volumes do not support [Dynamic Volume Provisioning](#). Options that allow combining Dynamic Volume Provisioning with Local Persistent Volumes are provided by [OpenEBS](#). Particularly, [OpenEBS Local PV Hostpath](#) allows creating Kubernetes Local Persistent Volumes using a directory (Hostpath) on the node. Such volume can be further accessed by applications via [Storage Class](#) and [PersistentVolumeClaim](#).

Using it involves the following steps.

1. Install OpenEBS on your system along with the official [installation guide](#).
2. Define a new [Kubernetes Storage Class](#) with OpenEBS with the YAML file (e. g. `local-hostpath.yaml`) as follows:

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: localpv
  annotations:
    openebs.io/cas-type: local
    cas.openebs.io/config: |
      - name: StorageType
        value: hostpath
      - name: BasePath
        value: /var/local-hostpath
provisioner: openebs.io/local
reclaimPolicy: Delete
volumeBindingMode: WaitForFirstConsumer
```

Two things to edit in this example are the `metadata.name` key (you will use it as a storage class name) and the `value` option under the `cas.openebs.io/config` (it should point to an already existing directory on the local filesystem of your node).

When ready, apply the file with the `kubectl apply -f local-hostpath.yaml` command.

12.6 . Now you can deploy the Operator and Percona XtraDB Cluster using this

StorageClass in `deploy/cr.yaml` :

```
...
volumeSpec:
  persistentVolumeClaim:
    storageClassName: localpv
    accessModes: [ "ReadWriteOnce" ]
  resources:
    requests:
      storage: 200Gi
```

Note

There are other storage options provided by the OpenEBS, which may be helpful within your cluster setup. Look at the [OpenEBS for the Management of Kubernetes Storage Volumes](#) blog post for more examples. Also, consider looking at the [Measuring OpenEBS Local Volume Performance Overhead in Kubernetes](#) post.

Last update: 2022-09-22

12.7 Define environment variables

Sometimes you need to define new environment variables to provide additional configuration for the components of your cluster. For example, you can use it to customize the configuration of HAProxy, or to add additional options for PMM Client.

The Operator can store environment variables in [Kubernetes Secrets](#). Here is an example with several HAProxy options:

```
apiVersion: v1
kind: Secret
metadata:
  name: my-env-var-secrets
type: Opaque
data:
  HA_CONNECTION_TIMEOUT: MTAwMA==
  OK_IF_DONOR: MQ==
  HA_SERVER_OPTIONS: Y2hLY2sgalW50ZXIgmZAwMDAgcmLzZSAXIGZhbGwgNSB3ZWlnahQgMQ==
```

As you can see, environment variables are stored as `data` - i.e., base64-encoded strings, so you'll need to encode the value of each variable. For example, To have `HA_CONNECTION_TIMEOUT` variable equal to `1000`, you can run `echo -n "1000" | base64 --wrap=0` (or just `echo -n "1000" | base64` in case of Apple macOS) in your local shell and get `MTAwMA==`.

Note

Similarly, you can read the list of options from a Base64-encoded string:

```
$ echo "MTAwMA==" | base64 --decode
```

When ready, apply the YAML file with the following command:

```
$ kubectl create -f deploy/my-env-secret.yaml
```

Put the name of this Secret to the `envVarsSecret` key either in `pxc`, `haproxy` or `proxysql` section of the `deploy/cr.yaml` configuration file:

```
haproxy:
  ....
  envVarsSecret: my-env-var-secrets
  ....
```

Now apply the `deploy/cr.yaml` file with the following command:

```
$ kubectl apply -f deploy/cr.yaml
```

Another example shows how to pass `LD_PRELOAD` environment variable with the alternative memory allocator library name to `mysqld`. It's often a recommended practice to try using an alternative allocator library for `mysqld` in case the memory usage is suspected to be higher than expected, and you can use `jemalloc` allocator already present in Percona XtraDB Cluster Pods with the following environment variable:

```
LD_PRELOAD=/usr/lib64/libjemalloc.so.1
```

Create a new YAML file with the contents similar to the previous example, but with `LD_PRELOAD` variable, stored as base64-encoded strings:

```
apiVersion: v1
kind: Secret
metadata:
  name: my-new-env-var-secrets
type: Opaque
data:
  LD_PRELOAD: L3Vzci9saWI2NC9saWJqZW1hbGxvYy5zby4x
```

If this YAML file was named `deploy/my-new-env-var-secret`, the command to apply it will be the following one:

```
$ kubectl create -f deploy/my-new-env-secret.yaml
```

Now put the name of this new Secret to the `envVarsSecret` key in `pxc` section of the `deploy/cr.yaml` configuration file:

```
pxc:
  ....
  envVarsSecret: my-new-env-var-secrets
  ....
```

Don't forget to apply the `deploy/cr.yaml` file, as usual:

```
$ kubectl apply -f deploy/cr.yaml
```

Last update: 2022-12-04

12.8 Configuring Load Balancing with HAProxy

Percona Operator for MySQL based on Percona XtraDB Cluster provides a choice of two cluster components to provide load balancing and proxy service: you can use either [HAProxy](#) or [ProxySQL](#). You can control which one to use, if any, by enabling or disabling via the `haproxy.enabled` and `proxysql.enabled` options in the `deploy/cr.yaml` configuration file.

Use the following command to enable HAProxy:

```
$ kubectl patch pxc cluster1 --type=merge --patch '{
  "spec": {
    "haproxy": {
      "enabled": true,
      "size": 3,
      "image": "percona/percona-xtradb-cluster-operator:1.12.0-haproxy" },
    "proxysql": { "enabled": false }
  }
}'
```

Note

For obvious reasons the Operator will not allow the simultaneous enabling of both HAProxy and ProxySQL.

The resulting HAProxy setup will contain two services:

- `cluster1-haproxy` service listening on ports 3306 (MySQL) and 3309 (the [proxy protocol](#)). This service is pointing to the number zero Percona XtraDB Cluster member (`cluster1-pxc-0`) by default when this member is available. If a zero member is not available, members are selected in descending order of their numbers (e.g. `cluster1-pxc-2`, then `cluster1-pxc-1`, etc.). This service can be used for both read and write load, or it can also be used just for write load (single writer mode) in setups with split write and read loads.
- `cluster1-haproxy-replicas` listening on port 3306 (MySQL). This service selects Percona XtraDB Cluster members to serve queries following the Round Robin load balancing algorithm.

When the cluster with HAProxy is upgraded, the following steps take place. First, reader members are upgraded one by one: the Operator waits until the upgraded Percona XtraDB Cluster member becomes synced, and then proceeds to upgrade the next member. When the upgrade is finished for all the readers, then the writer Percona XtraDB Cluster member is finally upgraded.

12.8.1 Passing custom configuration options to HAProxy

You can pass custom configuration to HAProxy in one of the following ways:

- edit the `deploy/cr.yaml` file,
- use a ConfigMap,
- use a Secret object.

Note

If you specify a custom HAProxy configuration in this way, the Operator doesn't provide its own HAProxy configuration file. That's why you should specify either a full set of configuration options or nothing.

Edit the `deploy/cr.yaml` file

You can add options from the `haproxy.cfg` configuration file by editing `haproxy.configuration` key in the `deploy/cr.yaml` file. Here is an example:

```
...
haproxy:
  enabled: true
  size: 3
  image: percona/percona-xtradb-cluster-operator:1.12.0-haproxy
  configuration: |
```

```

global
  maxconn 2048
  external-check
  stats socket /var/run/haproxy.sock mode 600 expose-fd listeners level user
defaults
  log global
  mode tcp
  retries 10
  timeout client 10000
  timeout connect 100500
  timeout server 10000
frontend galera-in
  bind *:3309 accept-proxy
  bind *:3306
  mode tcp
  option clitcpka
  default_backend galera-nodes
frontend galera-replica-in
  bind *:3309 accept-proxy
  bind *:3307
  mode tcp
  option clitcpka
  default_backend galera-replica-nodes

```

Use a ConfigMap

You can use a configmap and the cluster restart to reset configuration options. A configmap allows Kubernetes to pass or update configuration data inside a containerized application.

Use the `kubectl` command to create the configmap from external resources, for more information see [Configure a Pod to use a ConfigMap](#).

For example, you define a `haproxy.cfg` configuration file with the following setting:

```

global
  maxconn 2048
  external-check
  stats socket /var/run/haproxy.sock mode 600 expose-fd listeners level user
defaults
  log global
  mode tcp
  retries 10
  timeout client 10000
  timeout connect 100500
  timeout server 10000
frontend galera-in
  bind *:3309 accept-proxy
  bind *:3306
  mode tcp
  option clitcpka
  default_backend galera-nodes
frontend galera-replica-in
  bind *:3309 accept-proxy
  bind *:3307
  mode tcp
  option clitcpka
  default_backend galera-replica-nodes

```

You can create a configmap from the `haproxy.cfg` file with the `kubectl create configmap` command.

You should use the combination of the cluster name with the `-haproxy` suffix as the naming convention for the configmap. To find the cluster name, you can use the following command:

```
$ kubectl get pxc
```

The syntax for `kubectl create configmap` command is:

```
kubectl create configmap <cluster-name>-haproxy <resource-type=resource-name>
```

The following example defines `cluster1-haproxy` as the configmap name and the `haproxy.cfg` file as the data source:

```
$ kubectl create configmap cluster1-haproxy --from-file=haproxy.cfg
```

To view the created configmap, use the following command:

```
$ kubectl describe configmaps cluster1-haproxy
```

Use a Secret Object

The Operator can also store configuration options in [Kubernetes Secrets](#). This can be useful if you need additional protection for some sensitive data.

You should create a Secret object with a specific name, composed of your cluster name and the `haproxy` suffix.

Note

To find the cluster name, you can use the following command:

```
$ kubectl get pxc
```

Configuration options should be put inside a specific key inside of the `data` section. The name of this key is `haproxy.cfg` for ProxySQL Pods.

Actual options should be encoded with [Base64](#).

For example, let's define a `haproxy.cfg` configuration file and put there options we used in the previous example:

```
global
  maxconn 2048
  external-check
  stats socket /var/run/haproxy.sock mode 600 expose-fd listeners level user
defaults
  log global
  mode tcp
  retries 10
  timeout client 10000
  timeout connect 100500
  timeout server 10000
frontend galera-in
  bind *:3309 accept-proxy
  bind *:3306
  mode tcp
  option clitcpka
  default_backend galera-nodes
frontend galera-replica-in
  bind *:3309 accept-proxy
  bind *:3307
  mode tcp
  option clitcpka
  default_backend galera-replica-nodes
```

You can get a Base64 encoded string from your options via the command line as follows:

in Linux

```
$ cat haproxy.cfg | base64 --wrap=0
```

in macOS

```
$ cat haproxy.cfg | base64
```

Note

Similarly, you can read the list of options from a Base64 encoded string:

```
$ echo "IGdsb2JhbAogICBtYXhjb25uIDIwMDgKICAgZXh0ZXJlY2Y2sKICAgc3RhdHMgc29ja2V0\
IC92YXVlcnVUL2hhcHJveHkuc29jayBtb2RlIDYwM0CBLeHBvc2UetmQgbGZldGVuZXJzIGxldmV\
IHVzZXIKIGRlZmF1bHRzCiAgIGxvZyBnbG9iYmVwKICAgbW9kZSB0Y3AKICAgcmV0cmLlcyAxMAog\
ICB0aW1Lb3V0IGNsYWVudCAxMDAwAogICB0aW1Lb3V0IGNvbms1Y3QgMTAwNTAwCiAgIHRpbWV\
dXQgc2VydmlvYDEwMDAwCiBmcm9udGVuZCBnYXkLcmEtaW4KICAgYmLuZCAQ0jMzMDkgYmNjZXB0\
LXB3b3h5CiAgIGJpbmQgKj0zMzA2CiAgIG1vZGUgdG9wCiAgIG9udGlvbiBjbG9vZ3BrYQogICBk\
ZwZhdWx0X2JhY2tLbW9kZ2ZfSzxJhLW5vZGVzCiBmcm9udGVuZCBnYXkLcmEtcmlvY291bGog\
ICBiaW5kIC06MzMoOSBhY2NlcHQtcHJveHkKICAgYmLuZCAQ0jMzMDkKICAgbW9kZSB0Y3AKICAg\
b3B0aW9uIGNsXRjcGthCiAgIGRlZmF1bHRfYmFja2VuZCBnYXkLcmEtcmlvY291bGogLjY5Iub2Rlcwo=" | base64 --decode
```

Finally, use a yml file to create the Secret object. For example, you can create a `deploy/my-haproxy-secret.yml` file with the following contents:

```
apiVersion: v1
kind: Secret
```



```

metadata:
  name: cluster1-haproxy
data:
  my.cnf: "IGdsb2JhbAogICBtYXhjb25uIDIwNDgKICAgZkh0ZXJyYWRtY2hLY2sKICAgc3RhdHMgc29ja2V0\
IC92YXlvcnVul2hhcHJveHkuC29jayBtb2RlIDVwM0CBLeHBvc2UzZmQgbGZldGVuZjZlIGxldmVs\
IHVzZXIKIGRlZmF1bHRzCiAgIGxvZyBnbG9iYmVwKICAgbW9kZSB0Y3AKICAgcmV0cmllcyAxMAog\
ICB0aW1lb3V0IGNsaWVudCAxMDAwMAogICB0aW1lb3V0IGNvbms1Y3QgMTAwNTAwCiAgIHRpbWVv\
dXQgc2VydmlvYDEwMDAwCiBmcm9udGVuZCBnYXxLcmEtdW4KICAgYmLuZCAQ0jMzMDkgYWVjZXB0\
LXB3b3h5CiAgIGUjbWQgKjjozMzA2CiAgIG1vZG9gdG9wCiAgIG9wdGlvbiBjbGLOV3BrYQogICBk\
ZWZhdWx0X2JhY2t1bWQgZ2FsZXJhLW5vZGVzCiBmcm9udGVuZCBnYXxLcmEtdWVwLjYS1pbGog\
ICBiaW5kICo6MzMwOSBhY2NlCHQtcHJveHkKICAgYmLuZCAQ0jMzMDcKICAgbW9kZSB0Y3AKICAg\
b3B0aW9uIGNsaXRjcGthCiAgIGRlZmF1bHRfYmFja2VuZCBnYXxLcmEtdWVwLjYS1ub2Rlcwo="

```

When ready, apply it with the following command:

```
$ kubectl create -f deploy/my-haproxy-secret.yaml
```

Note

Do not forget to restart Percona XtraDB Cluster to ensure the cluster has updated the configuration.

12.8.2 Enabling the Proxy protocol

The Proxy protocol [allows](#) HAProxy to provide a real client address to Percona XtraDB Cluster.

Note

To use this feature, you should have a Percona XtraDB Cluster image version `8.0.21` or newer.

Normally Proxy protocol is disabled, and Percona XtraDB Cluster sees the IP address of the proxying server (HAProxy) instead of the real client address. But there are scenarios when making real client IP-address visible for Percona XtraDB Cluster is important: e.g. it allows to have privilege grants based on client/application address, and significantly enhance auditing.

You can enable Proxy protocol on Percona XtraDB Cluster by adding `proxy_protocol_networks` option to `pxc.configuration` key in the `deploy/cr.yaml` configuration file.

Note

Depending on the load balancer of your cloud provider, you may also need setting `haproxy.externaltrafficpolicy` option in `deploy/cr.yaml`.

More information about Proxy protocol can be found in the [official HAProxy documentation](#).

Last update: 2022-12-04

12.9 Configuring Load Balancing with ProxySQL

Percona Operator for MySQL based on Percona XtraDB Cluster provides a choice of two cluster components to provide load balancing and proxy service: you can use either [HAProxy](#) or [ProxySQL](#). You can control which one to use, if any, by enabling or disabling via the `haproxy.enabled` and `proxysql.enabled` options in the `deploy/cr.yaml` configuration file.

Use the following command to enable ProxySQL:

```
$ kubectl patch pxc cluster1 --type=merge --patch '{
"spec": {
  "proxysql": {
    "enabled": true,
    "size": 3,
    "image": "percona/percona-xtradb-cluster-operator:1.12.0-proxysql" },
  "haproxy": { "enabled": false }
}}
```

Note

For obvious reasons the Operator will not allow the simultaneous enabling of both HAProxy and ProxySQL.

The resulting setup will use the number zero Percona XtraDB Cluster member (`cluster1-pxc-0` by default) as writer.

When a cluster with ProxySQL is upgraded, the following steps take place. First, reader members are upgraded one by one: the Operator waits until the upgraded member shows up in ProxySQL with online status, and then proceeds to upgrade the next member. When the upgrade is finished for all the readers, then the writer Percona XtraDB Cluster member is finally upgraded.

Note

when both ProxySQL and Percona XtraDB Cluster are upgraded, they are upgraded in parallel.

12.9.1 Passing custom configuration options to ProxySQL

You can pass custom configuration to ProxySQL

- edit the `deploy/cr.yaml` file,
- use a ConfigMap,
- use a Secret object.

Note

If you specify a custom ProxySQL configuration in this way, ProxySQL will try to merge the passed parameters with the previously set configuration parameters, if any. If ProxySQL fails to merge some option, you will see a warning in its log.

Edit the `deploy/cr.yaml` file

You can add options from the `proxysql.cnf` configuration file by editing the `proxysql.configuration` key in the `deploy/cr.yaml` file. Here is an example:

```
...
proxysql:
  enabled: false
  size: 3
  image: percona/percona-xtradb-cluster-operator:1.12.0-proxysql
  configuration: |
    datadir="/var/lib/proxysql"

  admin_variables =
  {
    admin_credentials="proxysqladmin:admin_password"
```

```

mysql_ifaces="0.0.0.0:6032"
refresh_interval=2000

cluster_username="proxyadmin"
cluster_password="admin_password"
cluster_check_interval_ms=200
cluster_check_status_frequency=100
cluster_mysql_query_rules_save_to_disk=true
cluster_mysql_servers_save_to_disk=true
cluster_mysql_users_save_to_disk=true
cluster_proxysql_servers_save_to_disk=true
cluster_mysql_query_rules_diffs_before_sync=1
cluster_mysql_servers_diffs_before_sync=1
cluster_mysql_users_diffs_before_sync=1
cluster_proxysql_servers_diffs_before_sync=1
}

mysql_variables=
{
monitor_password="monitor"
monitor_galera_healthcheck_interval=1000
threads=2
max_connections=2048
default_query_delay=0
default_query_timeout=10000
poll_timeout=2000
interfaces="0.0.0.0:3306"
default_schema="information_schema"
stacksize=1048576
connect_timeout_server=10000
monitor_history=60000
monitor_connect_interval=20000
monitor_ping_interval=10000
ping_timeout_server=200
commands_stats=true
sessions_sort=true
have_ssl=true
ssl_p2s_ca="/etc/proxysql/ssl-internal/ca.crt"
ssl_p2s_cert="/etc/proxysql/ssl-internal/tls.crt"
ssl_p2s_key="/etc/proxysql/ssl-internal/tls.key"
ssl_p2s_cipher="ECDHE-RSA-AES128-GCM-SHA256"
}

```

Use a ConfigMap

You can use a configmap and the cluster restart to reset configuration options. A configmap allows Kubernetes to pass or update configuration data inside a containerized application.

Use the `kubectl` command to create the configmap from external resources, for more information see [Configure a Pod to use a ConfigMap](#).

For example, you define a `proxysql.cnf` configuration file with the following setting:

```

datadir="/var/lib/proxysql"

admin_variables =
{
admin_credentials="proxyadmin:admin_password"
mysql_ifaces="0.0.0.0:6032"
refresh_interval=2000

cluster_username="proxyadmin"
cluster_password="admin_password"
cluster_check_interval_ms=200
cluster_check_status_frequency=100
cluster_mysql_query_rules_save_to_disk=true
cluster_mysql_servers_save_to_disk=true
cluster_mysql_users_save_to_disk=true
cluster_proxysql_servers_save_to_disk=true
cluster_mysql_query_rules_diffs_before_sync=1
cluster_mysql_servers_diffs_before_sync=1
cluster_mysql_users_diffs_before_sync=1
cluster_proxysql_servers_diffs_before_sync=1
}

mysql_variables=
{
monitor_password="monitor"
monitor_galera_healthcheck_interval=1000
threads=2
max_connections=2048
default_query_delay=0
default_query_timeout=10000
poll_timeout=2000
interfaces="0.0.0.0:3306"
default_schema="information_schema"
stacksize=1048576
connect_timeout_server=10000
}

```

```

monitor_history=60000
monitor_connect_interval=20000
monitor_ping_interval=10000
ping_timeout_server=200
commands_stats=true
sessions_sort=true
have_ssl=true
ssl_p2s_ca="/etc/proxysql/ssl-internal/ca.crt"
ssl_p2s_cert="/etc/proxysql/ssl-internal/tls.crt"
ssl_p2s_key="/etc/proxysql/ssl-internal/tls.key"
ssl_p2s_cipher="ECDHE-RSA-AES128-GCM-SHA256"
}

```

You can create a configmap from the `proxysql.cnf` file with the `kubectl create configmap` command.

You should use the combination of the cluster name with the `-proxysql` suffix as the naming convention for the configmap. To find the cluster name, you can use the following command:

```
$ kubectl get pxc
```

The syntax for `kubectl create configmap` command is:

```
$ kubectl create configmap <cluster-name>-proxysql <resource-type=resource-name>
```

The following example defines `cluster1-proxysql` as the configmap name and the `proxysql.cnf` file as the data source:

```
$ kubectl create configmap cluster1-proxysql --from-file=proxysql.cnf
```

To view the created configmap, use the following command:

```
$ kubectl describe configmaps cluster1-proxysql
```

Use a Secret Object

The Operator can also store configuration options in [Kubernetes Secrets](#). This can be useful if you need additional protection for some sensitive data.

You should create a Secret object with a specific name, composed of your cluster name and the `proxysql` suffix.

Note

To find the cluster name, you can use the following command:

```
$ kubectl get pxc
```

Configuration options should be put inside a specific key inside of the `data` section. The name of this key is `proxysql.cnf` for ProxySQL Pods.

Actual options should be encoded with [Base64](#).

For example, let's define a `proxysql.cnf` configuration file and put there options we used in the previous example:

```

datadir="/var/lib/proxysql"

admin_variables =
{
  admin_credentials="proxyadmin:admin_password"
  mysql_ifaces="0.0.0.0:6032"
  refresh_interval=2000

  cluster_username="proxyadmin"
  cluster_password="admin_password"
  cluster_check_interval_ms=200
  cluster_check_status_frequency=100
  cluster_mysql_query_rules_save_to_disk=true
  cluster_mysql_servers_save_to_disk=true
  cluster_mysql_users_save_to_disk=true
  cluster_proxysql_servers_save_to_disk=true
  cluster_mysql_query_rules_diffs_before_sync=1
  cluster_mysql_servers_diffs_before_sync=1
  cluster_mysql_users_diffs_before_sync=1
  cluster_proxysql_servers_diffs_before_sync=1
}

```



```
Cn0KcM15c3FsX3ZhcmlhYmxiLz0KewogbW9uaXRvcL9wYXNzd29yZD0ibW9uaXRvcIiKIG1vbmL0\
b3JfZ2ZsZXJhX2hLYWx0aGNoZW50Lz0KewogbW9uaXRvcL9wYXNzd29yZD0ibW9uaXRvcIiKIG1vbmL0\
Y3Rpb25zPTIwNDgKIGRlZmF1bHRfcXVlcnLmFzZGVsYXk9MAogZGVmYXVsdF9xdWVyeV90aW1lb3V0\
PTEwMDAwCiBwb2x3RpbWVvdXQ9HjAwMAogalW50ZXJmYWNlcz0iMC4wLjAuMDozMzA2IggogZGVm\
YXVsdF9zY2hLbWE9ImLuZm9ybmF0aW9uX3NjaGVtYSIKIHNOYWwrc2L6ZT0xMDQ4NTc2CiBjb25u\
ZWNOX3RpbWVvdXRfc2VydmlvPTAwMDAwCiBtb25pdG9yX2hpc3Rvcnk9jAwMDAKIG1vbmL0b3Jf\
Y29ubmVjdF9pbmRlcnZhbD0yMDAwMAogbW9uaXRvcL9wYXNzd29yZD0ibW9uaXRvcIiBwa5n\
X3RpbWVvdXRfc2VydmlvPTAwMDAwY29tbWVzdHdHdHJlZQogc2Vzc2LvbWVzc29ydD10\
cnVlCiBoYXZlX3Nzbd10cnVlCiBzc2xfcDZjX2NhPSVvZXRjL3Byb3h5c3FsL3NzbC1pbmRlcm5h\
bC9jY55jcnQiCiBzc2xfcDZjX2NlcnQ9Ii9LdG9ybmVhdHdHdHJlZQogc2Vzc2LvbWVzc29ydD10\
cnVlCiBzc2xfcDZjX2tLeT0iL2V0Yy9wcm94eXNxbC9zc2wtaw50ZkxwYVwvdGxzLmtLeSikIHnz\
bF9wbnNfY2lwaGVyPSJFQ0RIRS1SU0EtQUVtTMTI4LUdDTs1TSEEyNTYiCn0K"
```

When ready, apply it with the following command:

```
$ kubectl create -f deploy/my-proxySQL-secret.yaml
```

Note

Do not forget to restart Percona XtraDB Cluster to ensure the cluster has updated the configuration.

12.9.2 Accessing the ProxySQL Admin Interface

You can use [ProxySQL admin interface](#) to configure its settings.

Configuring ProxySQL in this way means connecting to it using the MySQL protocol, and two things are needed to do it:

- the ProxySQL Pod name
- the ProxySQL admin password

You can find out ProxySQL Pod name with the `kubectl get pods` command, which will have the following output:

```
$ kubectl get pods
NAME                                READY   STATUS    RESTARTS   AGE
cluster1-pxc-node-0                 1/1     Running   0           5m
cluster1-pxc-node-1                 1/1     Running   0           4m
cluster1-pxc-node-2                 1/1     Running   0           2m
cluster1-proxySQL-0                 1/1     Running   0           5m
percona-xtradb-cluster-operator-dc67778fd-qtspz 1/1     Running   0           6m
```

The next command will print you the needed admin password:

```
$ kubectl get secrets $(kubectl get pxc -o jsonpath='{.items[0].spec.secretsName}') -o template='{{ .data.proxyadmin | base64decode }}'
```

When both Pod name and admin password are known, connect to the ProxySQL as follows, substituting `cluster1-proxySQL-0` with the actual Pod name and `admin_password` with the actual password:

```
$ kubectl exec -it cluster1-proxySQL-0 -- mysql -h127.0.0.1 -P6032 -uproxyadmin -padmin_password
```

Last update: 2022-12-04

12.10 Transport Layer Security (TLS)

The Percona Operator for MySQL uses Transport Layer Security (TLS) cryptographic protocol for the following types of communication:

- Internal - communication between Percona XtraDB Cluster instances,
- External - communication between the client application and ProxySQL.

The internal certificate is also used as an authorization method.

TLS security can be configured in several ways. By default, the Operator generates long-term certificates automatically if there are no certificate secrets available. Other options are the following ones:

- The Operator can use a specifically installed *cert-manager*, which will automatically generate and renew short-term TLS certificates,
- Certificates can be generated manually.

You can also use pre-generated certificates available in the `deploy/ssl-secrets.yaml` file for test purposes, but we strongly recommend avoiding their usage on any production system!

The following subsections explain how to configure TLS security with the Operator yourself, as well as how to temporarily disable it if needed.

12.10.1 Install and use the *cert-manager*

About the *cert-manager*

A [cert-manager](#) is a Kubernetes certificate management controller which is widely used to automate the management and issuance of TLS certificates. It is community-driven, and open source.

When you have already installed *cert-manager* and deploy the operator, the operator requests a certificate from the *cert-manager*. The *cert-manager* acts as a self-signed issuer and generates certificates. The Percona Operator self-signed issuer is local to the operator namespace. This self-signed issuer is created because Percona XtraDB Cluster requires all certificates issued by the same .

Self-signed issuer allows you to deploy and use the Percona Operator without creating a clusterissuer separately.

Installation of the *cert-manager*

The steps to install the *cert-manager* are the following:

- Create a namespace,
- Disable resource validations on the cert-manager namespace,
- Install the cert-manager.

The following commands perform all the needed actions:

```
$ kubectl create namespace cert-manager
$ kubectl label namespace cert-manager certmanager.k8s.io/disable-validation=true
$ kubectl apply -f https://github.com/jetstack/cert-manager/releases/download/v1.8.0/cert-manager.yaml
```

After the installation, you can verify the *cert-manager* by running the following command:

```
$ kubectl get pods -n cert-manager
```

The result should display the *cert-manager* and webhook active and running.

12.10.2 Generate certificates manually

To generate certificates manually, follow these steps:

1. Provision a Certificate Authority (CA) to generate TLS certificates
2. Generate a CA key and certificate file with the server details
3. Create the server TLS certificates using the CA keys, certs, and server details

The set of commands generate certificates with the following attributes:

- `Server.pem` - Certificate
- `Server-key.pem` - the private key
- `ca.pem` - Certificate Authority

You should generate certificates twice: one set is for external communications, and another set is for internal ones. A secret created for the external use must be added to `cr.yaml/spec/secretsName`. A certificate generated for internal communications must be added to the `cr.yaml/spec/sslInternalSecretName`.

```
$ cat <<EOF | cfssl gencert -initca - | cfssljson -bare ca
{
  "CN": "Root CA",
  "key": {
    "algo": "rsa",
    "size": 2048
  }
}
EOF

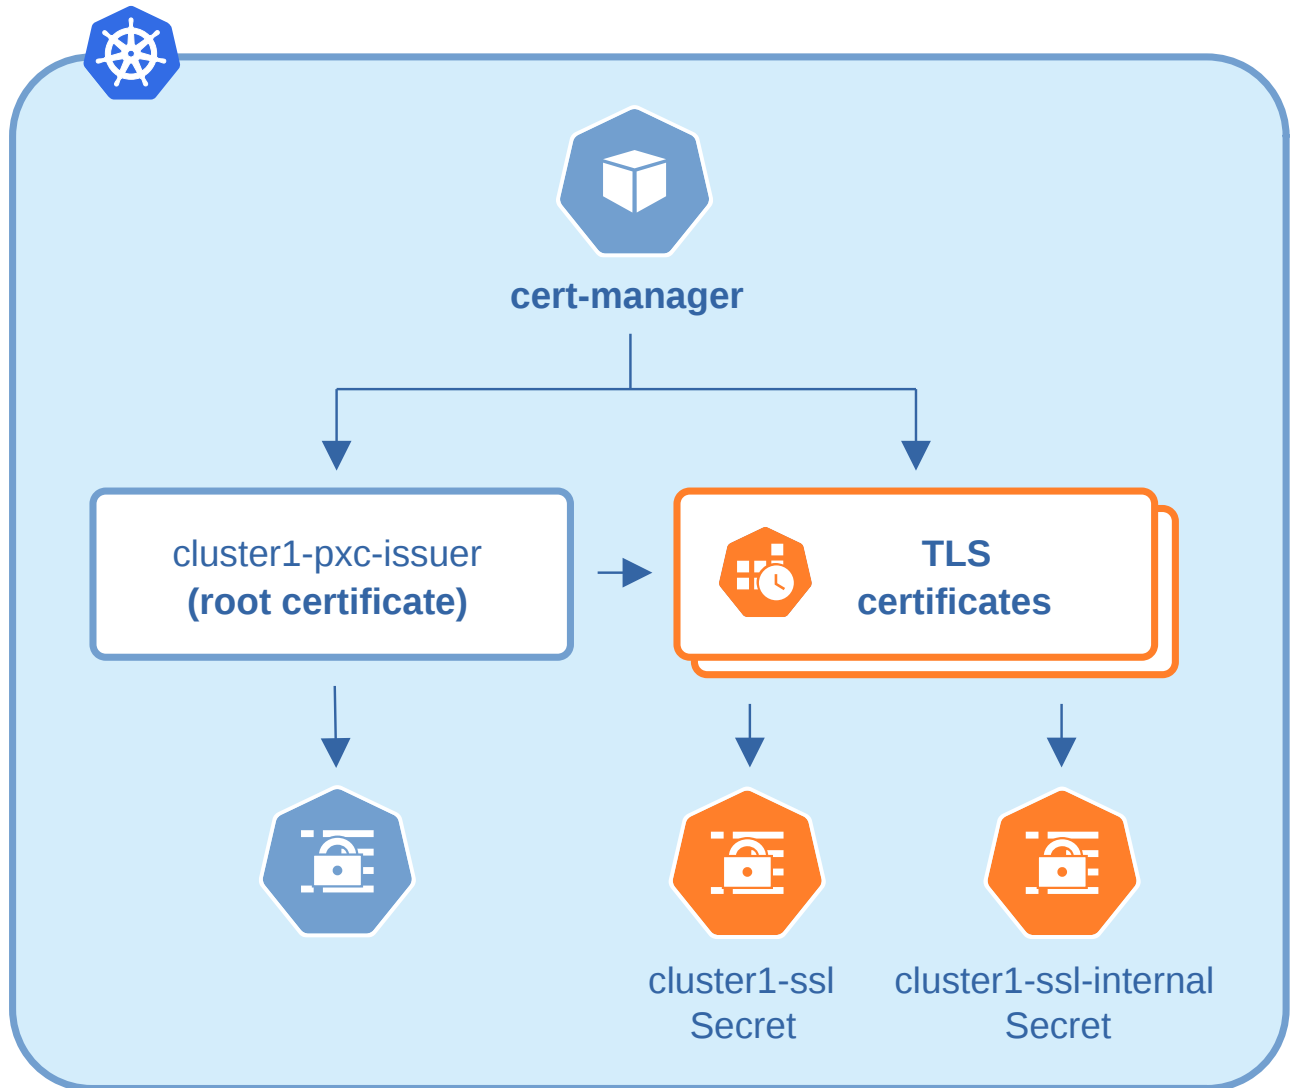
$ cat <<EOF | cfssl gencert -ca=ca.pem -ca-key=ca-key.pem - | cfssljson -bare server
{
  "hosts": [
    "${CLUSTER_NAME}-proxysql",
    "${CLUSTER_NAME}-proxysql-unready",
    "${CLUSTER_NAME}-pxc"
  ],
  "CN": "${CLUSTER_NAME}-pxc",
  "key": {
    "algo": "rsa",
    "size": 2048
  }
}
EOF

$ kubectl create secret generic cluster1-ssl --from-file=tls.crt=server.pem --
from-file=tls.key=server-key.pem --from-file=ca.crt=ca.pem --
type=kubernetes.io/tls
```

12.10.3 Update certificates

If a cert-manager is used, it should take care of updating the certificates. If you generate certificates manually, you should take care of updating them in proper time.

TLS certificates issued by cert-manager are short-term ones. Starting from the Operator version 1.9.0 cert-manager issues TLS certificates for 3 months, while root certificate is valid for 3 years. This allows to reissue TLS certificates automatically on schedule and without downtime.



Versions of the Operator prior 1.9.0 have used 3 month root certificate, which caused issues with the automatic TLS certificates update. If that's your case, you can make the Operator update along with the [official instruction](#).

Note

If you use the cert-manager version earlier than 1.9.0, and you would like to avoid downtime while updating the certificates after the Operator update to 1.9.0 or newer version, force the certificates regeneration by a cert-manager.

Check your certificates for expiration

1. First, check the necessary secrets names (`cluster1-ssl` and `cluster1-ssl-internal` by default):

```
$ kubectl get certificate
```

You will have the following response:

NAME	READY	SECRET	AGE
cluster1-ssl	True	cluster1-ssl	49m
cluster1-ssl-internal	True	cluster1-ssl-internal	49m

2. Optionally you can also check that the certificates issuer is up and running:

```
$ kubectl get issuer
```

The response should be as follows:

```
NAME          READY   AGE
cluster1-pxc-ca  True    49m
```

3. Now use the following command to find out the certificates validity dates, substituting Secrets names if necessary:

```
$ {
  kubectl get secret/cluster1-ssl-internal -o jsonpath='{.data.tls.crt}' | base64 --decode | openssl x509 -inform pem -noout -text | grep "Not After"
  kubectl get secret/cluster1-ssl -o jsonpath='{.data.ca.crt}' | base64 --decode | openssl x509 -inform pem -noout -text | grep "Not After"
}
```

The resulting output will be self-explanatory:

```
Not After : Sep 15 11:04:53 2021 GMT
Not After : Sep 15 11:04:53 2021 GMT
```

Update certificates without downtime

If you don't use cert-manager and have *created certificates manually*, you can follow the next steps to perform a no-downtime update of these certificates *if they are still valid*.

Note

For already expired certificates, follow the alternative way.

Having non-expired certificates, you can roll out new certificates (both CA and TLS) with the Operator as follows.

1. Generate a new CA certificate (`ca.pem`). Optionally you can also generate a new TLS certificate and a key for it, but those can be generated later on step 6.
2. Get the current CA (`ca.pem.old`) and TLS (`tls.pem.old`) certificates and the TLS certificate key (`tls.key.old`):

```
$ kubectl get secret/cluster1-ssl-internal -o jsonpath='{.data.ca\.crt}' | base64 --decode > ca.pem.old
$ kubectl get secret/cluster1-ssl-internal -o jsonpath='{.data.tls\.crt}' | base64 --decode > tls.pem.old
$ kubectl get secret/cluster1-ssl-internal -o jsonpath='{.data.tls\.key}' | base64 --decode > tls.key.old
```

3. Combine new and current `ca.pem` into a `ca.pem.combined` file:

```
$ cat ca.pem ca.pem.old >> ca.pem.combined
```

4. Create a new Secrets object with *old* TLS certificate (`tls.pem.old`) and key (`tls.key.old`), but a *new combined* `ca.pem` (`ca.pem.combined`):

```
$ kubectl delete secret/cluster1-ssl-internal
$ kubectl create secret generic cluster1-ssl-internal --from-file=tls.crt=tls.pem.old --from-file=tls.key=tls.key.old --from-file=ca.crt=ca.pem.combined --type=kubernetes.io/tls
```

5. The cluster will go through a rolling reconciliation, but it will do it without problems, as every node has old TLS certificate/key, and both new and old CA certificates.
6. If new TLS certificate and key weren't generated on step 1, do that now.
7. Create a new Secrets object for the second time: use new TLS certificate (`server.pem` in the example) and its key (`server-key.pem`), and again the combined CA certificate (`ca.pem.combined`):

```
$ kubectl delete secret/cluster1-ssl-internal
$ kubectl create secret generic cluster1-ssl-internal --from-file=tls.crt=server.pem --from-file=tls.key=server-key.pem --from-file=ca.crt=ca.pem.combined --type=kubernetes.io/tls
```

8. The cluster will go through a rolling reconciliation, but it will do it without problems, as every node already has a new CA certificate (as a part of the combined CA certificate), and can successfully allow joiners with new TLS certificate to join. Joiner node also has a combined CA certificate, so it can authenticate against older TLS certificate.
9. Create a final Secrets object: use new TLS certificate (`server.pmm`) and its key (`server-key.pem`), and just the new CA certificate (`ca.pem`):

```
$ kubectl delete secret/cluster1-ssl-internal
$ kubectl create secret generic cluster1-ssl-internal --from-file=tls.crt=server.pem --from-file=tls.key=server-key.pem --from-file=ca.crt=ca.pem --type=kubernetes.io/tls
```

10. The cluster will go through a rolling reconciliation, but it will do it without problems: the old CA certificate is removed, and every node is already using new TLS certificate and no nodes rely on the old CA certificate any more.

Update certificates with downtime

If your certificates have been already expired (or if you continue to use the Operator version prior to 1.9.0), you should move through the *pause - update Secrets - unpause* route as follows.

1. Pause the cluster [in a standard way](#), and make sure it has reached its paused state.
2. If cert-manager is used, delete issuer and TLS certificates:

```
{
  kubectl delete issuer/cluster1-pxc-ca
  kubectl delete certificate/cluster1-ssl certificate/cluster1-ssl-internal
}
```

3. Delete Secrets to force the SSL reconciliation:

```
$ kubectl delete secret/cluster1-ssl secret/cluster1-ssl-internal
```

4. Check certificates to make sure reconciliation have succeeded.
5. Unpause the cluster [in a standard way](#), and make sure it has reached its running state.

Keep certificates after deleting the cluster

In case of cluster deletion, objects, created for SSL (Secret, certificate, and issuer) are not deleted by default.

If the user wants the cleanup of objects created for SSL, there is a [finalizers.delete-ssl](#) option in `deploy/cr.yaml`: if this finalizer is set, the Operator will delete Secret, certificate and issuer after the cluster deletion event.

12.10.4 Run Percona XtraDB Cluster without TLS

Omitting TLS is also possible, but we recommend that you run your cluster with the TLS protocol enabled.

To disable TLS protocol (e.g. for demonstration purposes) edit the `cr.yaml/spec/allowUnsafeConfigurations` setting to `true` and make sure that there are no certificate secrets available.

Last update: 2022-12-07

12.11 Data at Rest Encryption

Full data at rest encryption in Percona XtraDB Cluster is supported by the Operator since version 1.4.0.

Note

Data at rest means inactive data stored as files, database records, etc.

To implement these features, the Operator uses `keyring_vault` plugin, which ships with Percona XtraDB Cluster, and utilizes [HashiCorp Vault](#) storage for encryption keys.

12.11.1 Installing Vault

The following steps will deploy Vault on Kubernetes with the [Helm 3 package manager](#). Other Vault installation methods should also work, so the instruction placed here is not obligatory and is for illustration purposes. Read more about installation in Vault's [documentation](#).

1. Add helm repo and install:

```
$ helm repo add hashicorp https://helm.releases.hashicorp.com
"hashicorp" has been added to your repositories

$ helm install vault hashicorp/vault
```

2. After the installation, Vault should be first initialized and then unsealed. Initializing Vault is done with the following commands:

```
$ kubectl exec -it pod/vault-service-0 -- vault operator init -key-shares=1 -key-threshold=1 -format=json > /tmp/vault-init
$ unsealKey=$(jq -r ".unseal_keys_b64[]" < /tmp/vault-init)
```

To unseal Vault, execute the following command **for each Pod** of Vault running:

```
$ kubectl exec -it pod/vault-service-0 -- vault operator unseal "$unsealKey"
```

12.11.2 Configuring Vault

1. First, you should enable secrets within Vault. For this you will need a [Vault token](#). Percona XtraDB Cluster can use any regular token which allows all operations inside the secrets mount point. In the following example we are using the `root token` to be sure the permissions requirement is met, but actually there is no need in root permissions. We don't recommend using the root token on the production system.

```
$ cat /tmp/vault-init | jq -r ".root_token"
```

The output will be like follows:

```
s.VgQvaXl8xGF01RUxAPbPbsfN
```

Now login to Vault with this token and enable the "pxc-secret" secrets path:

```
$ kubectl exec -it vault-service-0 -- /bin/sh
$ vault login s.VgQvaXl8xGF01RUxAPbPbsfN
$ vault secrets enable --version=1 -path=pxc-secret kv
```

Note

You can also enable audit, which is not mandatory, but useful:

```
$ vault audit enable file file_path=/vault/vault-audit.log
```

2. To enable Vault secret within Kubernetes, create and apply the YAML file, as described further.
 - a. To access the Vault server via HTTP, follow the next YAML file example:

```
apiVersion: v1
kind: Secret
metadata:
  name: some-name-vault
type: Opaque
stringData:
  keyring_vault.conf: |-
    token = s.VgQvaXl8xGF01RUxAPbPbsfN
    vault_url = vault-service.vault-service.svc.cluster.local
    secret_mount_point = pxc-secret
```

Note

the `name` key in the above file should be equal to the `spec.vaultSecretName` key from the `deploy/cr.yaml` configuration file.

- b. To turn on TLS and access the Vault server via HTTPS, you should do two more things:
 - add one more item to the secret: the contents of the `ca.cert` file with your certificate,
 - store the path to this file in the `vault_ca` key.

```
apiVersion: v1
kind: Secret
metadata:
  name: some-name-vault
type: Opaque
stringData:
  keyring_vault.conf: |-
    token = s.VgQvaXl8xGF01RUxAPbPbsfN
    vault_url = https://vault-service.vault-service.svc.cluster.local
    secret_mount_point = pxc-secret
    vault_ca = /etc/mysql/vault-keyring-secret/ca.cert
  ca.cert: |-
    -----BEGIN CERTIFICATE-----
    MIIEczCCAlugAwIBAgIBADANBgkqhkiG9w0BAQQFAD. .AkGA1UEBhMCR0Ix
    EzARBgNVBAGTCLNvbWUtU3RhdGUxFTAsBgNVBAoTCO. .0EgTHRkMTcwNQYD
    7vQMFXdGSRrXNGRgnX+vNDZ3/zWIOjOdTckNngEpVn. .HoX
    -----END CERTIFICATE-----
```

Note

the `name` key in the above file should be equal to the `spec.vaultSecretName` key from the `deploy/cr.yaml` configuration file.

Note

For technical reasons the `vault_ca` key should either exist or not exist in the YAML file; commented option like `#vault_ca = ...` is not acceptable.

More details on how to install and configure Vault can be found [in the official documentation](#).

12.11.3 Using the encryption

If using *Percona XtraDB Cluster 5.7*, you should turn encryption on explicitly when you create a table or a tablespace. This can be done by adding the `ENCRYPTION='Y'` part to your SQL statement, like in the following example:

```
CREATE TABLE t1 (c1 INT, PRIMARY KEY pk(c1)) ENCRYPTION='Y';
CREATE TABLESPACE foo ADD DATAFILE 'foo.ibd' ENCRYPTION='Y';
```

Note

See more details on encryption in Percona XtraDB Cluster 5.7 [here](#).

If using *Percona XtraDB Cluster 8.0*, the encryption is turned on by default (in case if Vault is configured).

The following table presents the default values of the [correspondent my.cnf configuration options](#):

Option	Default value
early-plugin-load	keyring_vault.so
keyring_vault_config	/etc/mysql/vault-keyring-secret/keyring_vault.conf
default_table_encryption	ON
table_encryption_privilege_check	ON
innodb_undo_log_encrypt	ON
innodb_redo_log_encrypt	ON
binlog_encryption	ON
binlog_rotate_encryption_master_key_at_startup	ON
innodb_temp_tablespace_encrypt	ON
innodb_parallel_dblwr_encrypt	ON
innodb_encrypt_online_alter_logs	ON
encrypt_tmp_files	ON

.....
Last update: 2022-12-04

12.12 Telemetry

The Telemetry function enables the Operator gathering and sending basic anonymous data to Percona, which helps us to determine where to focus the development and what is the uptake for each release of Operator.

The following information is gathered:

- ID of the Custom Resource (the `metadata.uid` field)
- Kubernetes version
- Platform (is it Kubernetes or Openshift)
- PMM Version
- Operator version
- Percona XtraDB Cluster version
- HAProxy version
- ProxySQL version
- Percona XtraBackup version
- Is Operator deployed in a cluster-wide mode

We do not gather anything that identify a system, but the following thing should be mentioned: Custom Resource ID is a unique ID generated by Kubernetes for each Custom Resource.

Telemetry is enabled by default and is sent to the [Version Service server](#) when the Operator connects to it at scheduled times to obtain fresh information about version numbers and valid image paths needed for the upgrade.

The landing page for this service, check.percona.com, explains what this service is.

You can disable telemetry with a special option when installing the Operator:

- if you [install the Operator with helm](#), use the following installation command:

```
$ helm install my-db percona/pxc-db --version 1.12.0 --namespace my-namespace --set disable_telemetry="true"
```

- if you don't use helm for installation, you have to edit the `operator.yaml` before applying it with the `kubectl apply -f deploy/operator.yaml` command. Open the `operator.yaml` file with your text editor, find the value of the `DISABLE_TELEMETRY` environment variable and set it to `true`:

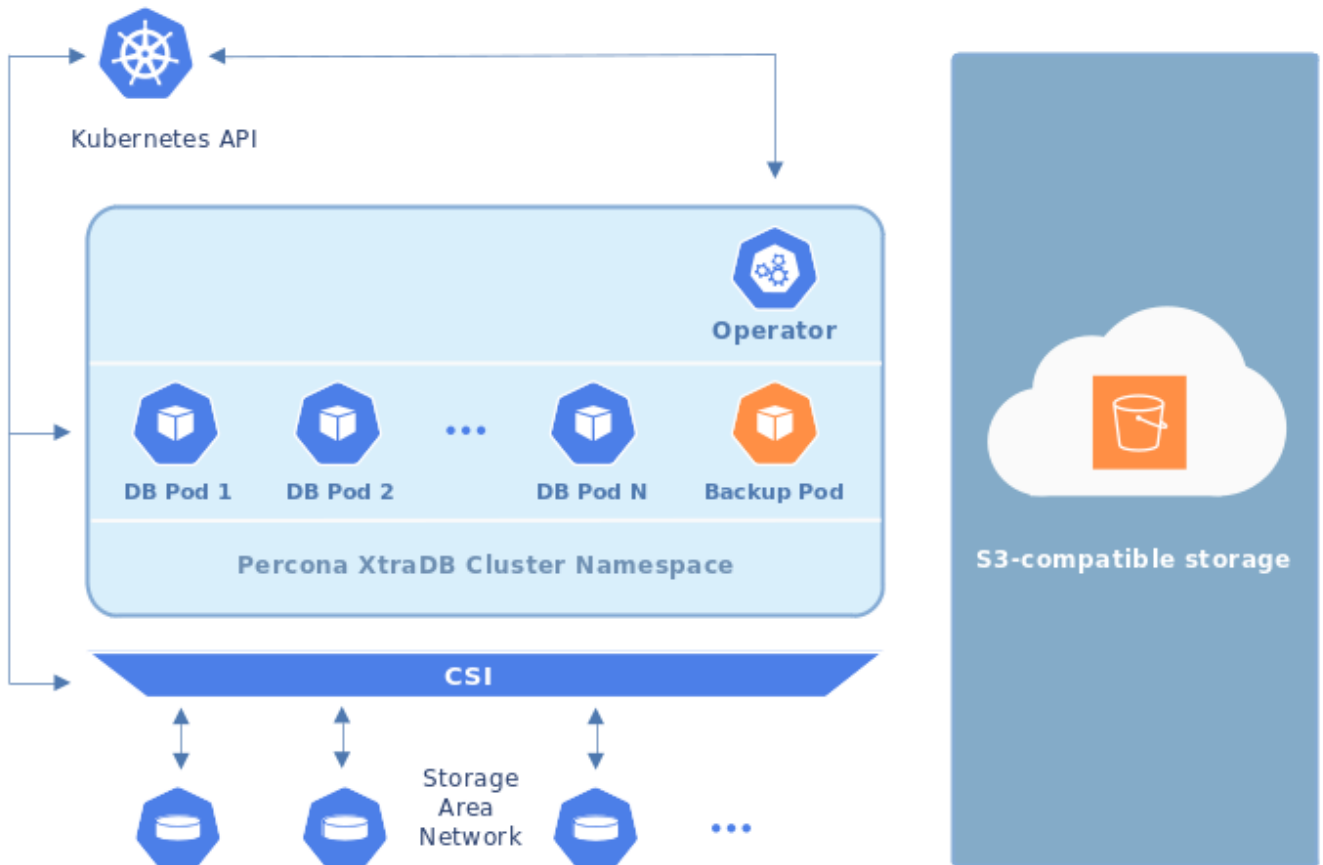
```
env:
  ...
  - name: DISABLE_TELEMETRY
    value: "true"
  ...
```

Last update: 2022-12-07

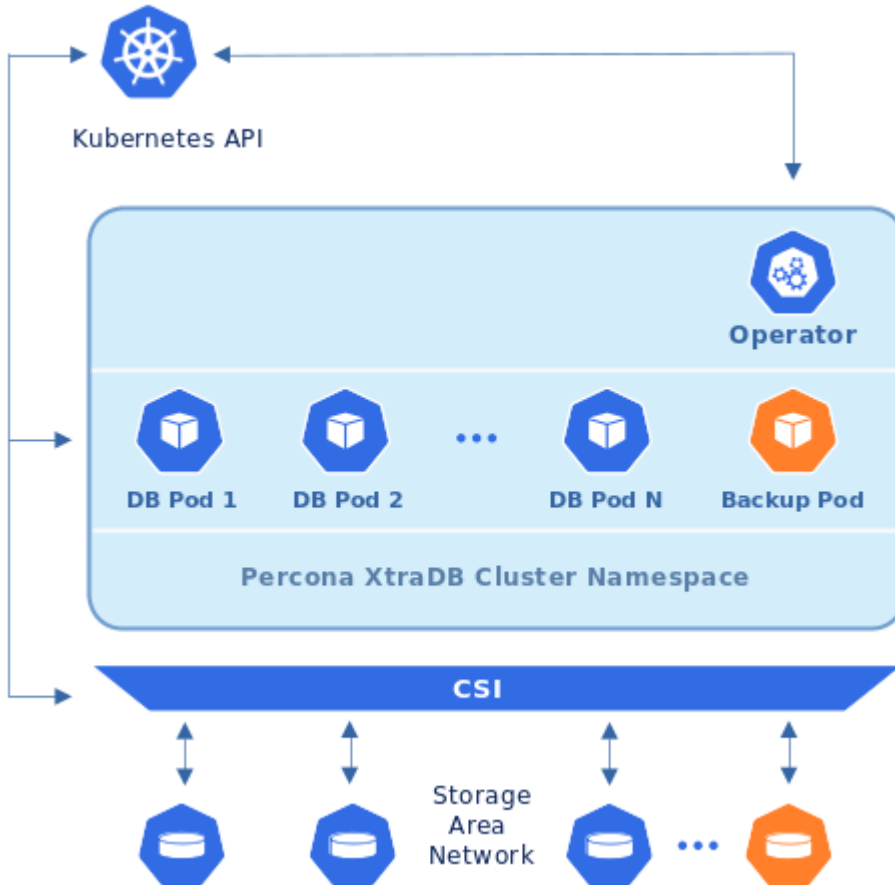
13. Management

13.1 Providing Backups

The Operator usually stores Percona XtraDB Cluster backups on [Amazon S3 or S3-compatible storage](#) outside the Kubernetes cluster: The Operator usually stores Percona XtraDB Cluster backups outside the Kubernetes cluster, on [Amazon S3 or S3-compatible storage](#), or on [Azure Blob Storage](#):



But storing backups on [Persistent Volumes](#) inside the Kubernetes cluster is also possible:



The Operator allows doing backups in two ways. *Scheduled backups* are configured in the `deploy/cr.yaml` file to be executed automatically in proper time. *On-demand backups* can be done manually at any moment.

13.1.1 Making scheduled backups

Backups schedule is defined in the `backup` section of the `deploy/cr.yaml` file. This section contains following subsections:

- `storages` subsection contains data and configuration needed to store backups,
- `schedule` subsection allows to actually schedule backups (the schedule is specified in crontab format).

Backups on Amazon S3 or S3-compatible storage

Since backups are stored separately on the Amazon S3, a secret with `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY` should be present on the Kubernetes cluster. The secrets file with these base64-encoded keys should be created: for example `deploy/backup-s3.yaml` file with the following contents:

```
apiVersion: v1
kind: Secret
metadata:
  name: my-cluster-name-backup-s3
type: Opaque
data:
  AWS_ACCESS_KEY_ID: UKVQTEFDRS1XSVRILUFUy1BQ0NFU1M6SOVZ
  AWS_SECRET_ACCESS_KEY: UKVQTEFDRS1XSVRILUFUy1TRUNSRVQ6SOVZ
```

Note

The following command can be used to get a base64-encoded string from a plain text one:

in Linux

```
$ echo -n 'plain-text-string' | base64 --wrap=0
```

in macOS

```
$ echo -n 'plain-text-string' | base64
```

The `name` value is the [Kubernetes secret](#) name which will be used further, and `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY` are the keys to access S3 storage (and obviously they should contain proper values to make this access possible). To have effect secrets file should be applied with the appropriate command to create the secret object, e.g. `kubectl apply -f deploy/backup-s3.yaml` (for Kubernetes).

Note

In case if the previous backup attempt fails (because of a temporary networking problem, etc.) the backup job tries to delete the unsuccessful backup leftovers first, and then makes a retry. Therefore there will be no backup retry without [DELETE permissions to the objects in the bucket](#). Also, setting [Google Cloud Storage Retention Period](#) can cause a similar problem.

All the data needed to access the S3-compatible cloud to store backups (credentials, name of the [bucket](#) to keep backups in, etc.) should be put into the `backup.storages` subsection, and `backup.schedule` subsection should actually schedule backups in crontab-compatible way. Here is an example of [deploy/cr.yaml](#) which uses Amazon S3 storage for backups:

```
...
backup:
  ...
  storages:
    s3-us-west:
      type: s3
      s3:
        bucket: S3-BACKUP-BUCKET-NAME-HERE
        region: us-west-2
        credentialsSecret: my-cluster-name-backup-s3
  ...
  schedule:
  - name: "sat-night-backup"
    schedule: "0 0 * * 6"
    keep: 3
    storageName: s3-us-west
  ...
```

if you use some S3-compatible storage instead of the original Amazon S3, the `endpointUrl` is needed in the `s3` subsection which points to the actual cloud used for backups and is specific to the cloud provider. For example, using [Google Cloud](#) involves the [following endpointUrl](#):

```
endpointUrl: https://storage.googleapis.com
```

The options within these three subsections are further explained in the [Custom Resource options](#).

One option which should be mentioned separately is `credentialsSecret` which is a [Kubernetes secret](#) for backups. Value of this key should be the same as the name used to create the secret object (`my-cluster-name-backup-s3` in the last example).

Backups on Microsoft Azure Blob storage

Since backups are stored separately on [Azure Blob Storage](#), a secret with `AZURE_STORAGE_ACCOUNT_NAME` and `AZURE_STORAGE_ACCOUNT_KEY` should be present on the Kubernetes cluster. The secrets file with these base64-encoded keys should be created: for example `deploy/backup-azure.yaml` file with the following contents.

```
apiVersion: v1
kind: Secret
metadata:
  name: my-cluster-azure-secret
type: Opaque
```

```
data:
  AZURE_STORAGE_ACCOUNT_NAME: UKVQTEFDRS1XSVRILUFXY1BQ0NFU1MESOVZ
  AZURE_STORAGE_ACCOUNT_KEY: UKVQTEFDRS1XSVRILUFXY1TRUNSRVQtSOVZ
```

Note

The following command can be used to get a base64-encoded string from a plain text one:

in Linux

```
$ echo -n 'plain-text-string' | base64 --wrap=0
```

in macOS

```
$ echo -n 'plain-text-string' | base64
```

The `name` value is the [Kubernetes secret](#) name which will be used further. The `AZURE_STORAGE_ACCOUNT_NAME` and `AZURE_STORAGE_ACCOUNT_KEY` credentials will be used to access the storage (and obviously they should contain proper values to make this access possible). To have effect, secrets file should be applied with the appropriate command to create the Secrets object, e.g. `kubectl apply -f deploy/backup-azure.yaml` (for Kubernetes).

All the data needed to access the Azure Blob storage to store backups (credentials, name of the [container](#) to keep backups in, etc.) should be put into the `backup.storages` subsection, and `backup.schedule` subsection should actually schedule backups in crontab-compatible way. Here is an example of [deploy/cr.yaml](#) which uses Azure Blob storage for backups:

```
...
backup:
  enabled: true
  ...
  storages:
    azure-blob:
      type: azure
      azure:
        credentialsSecret: my-cluster-azure-secret
        container: <your-container-name>
      ...
  schedule:
    - name: "sat-night-backup"
      schedule: "0 0 * * 6"
      keep: 3
      storageName: azure-blob
  ...
```

The options within these three subsections are further explained in the [Operator Custom Resource options](#).

One option which should be mentioned separately is `credentialsSecret` which is a [Kubernetes secret](#) for backups. Value of this key should be the same as the name used to create the secret object (`my-cluster-azure-secret` in the last example).

13.1.2 Making on-demand backup

To make an on-demand backup, the user should first configure the backup storage in the `backup.storages` subsection of the `deploy/cr.yaml` configuration file in a same way it was done for scheduled backups. When the `deploy/cr.yaml` file contains correctly configured storage and is applied with `kubectl` command, use *a special backup configuration YAML file* with the following contents:

- **backup name** in the `metadata.name` key,
- **Percona XtraDB Cluster name** in the `spec.pxcCluster` key,
- **storage name** from `deploy/cr.yaml` in the `spec.storageName` key,
- **S3 backup finalizer** set by the `metadata.finalizers.delete-s3-backup` key (it triggers the actual deletion of backup files from the S3 bucket when there is a manual or scheduled removal of the corresponding backup object).

The example of the backup configuration file is [deploy/backup/backup.yaml](#).

When the backup destination is configured and applied with `kubectl apply -f deploy/cr.yaml` command, the actual backup command is executed:

```
$ kubectl apply -f deploy/backup/backup.yaml
```

Note

Storing backup settings in a separate file can be replaced by passing its content to the `kubectl apply` command as follows:

```
$ cat <<EOF | kubectl apply -f-
apiVersion: pxc.percona.com/v1
kind: PerconaXtraDBClusterBackup
metadata:
  finalizers:
    - delete-s3-backup
  name: backup1
spec:
  pxcCluster: cluster1
  storageName: s3-us-west
EOF
```

13.1.3 Storing binary logs for point-in-time recovery

Point-in-time recovery functionality allows users to roll back the cluster to a specific transaction, time (or even skip a transaction in some cases). Technically, this feature involves continuously saving binary log updates to the backup storage. Point-in-time recovery is off by default and is supported by the Operator only with Percona XtraDB Cluster versions starting from 8.0.21-12.1.

To be used, it requires setting a number of keys in the `pitr` subsection under the `backup` section of the `deploy/cr.yaml` file:

- `enabled` key should be set to `true`,
- `storageName` key should point to the name of the storage already configured in the `storages` subsection

Note

Both binlog and full backup should use s3-compatible storage to make point-in-time recovery work!

- `timeBetweenUploads` key specifies the number of seconds between running the binlog uploader.

Following example shows how the `pitr` subsection looks like:

```
backup:
  ...
  pitr:
    enabled: true
    storageName: s3-us-west
    timeBetweenUploads: 60
```

Note

Point-in-time recovery will be done for binlogs without any cluster-based filtering. Therefore it is recommended to use a separate storage, bucket, or directory to store binlogs for the cluster. Also, it is recommended to have empty bucket/directory which holds binlogs (with no binlogs or files from previous attempts or other clusters) when you enable point-in-time recovery.

Note

Purging binlogs before they are transferred to backup storage will break point-in-time recovery.

13.1.4 Storing backup on Persistent Volume

Here is an example of the `deploy/cr.yaml` backup section fragment, which configures a private volume for filesystem-type storage:

```
...
backup:
  ...
  storages:
    fs-pvc:
      type: filesystem
      volume:
        persistentVolumeClaim:
```

```

accessModes: [ "ReadWriteOnce" ]
resources:
  requests:
    storage: 6Gi
...

```

Note

Please take into account that 6Gi storage size specified in this example may be insufficient for the real-life setups; consider using tens or hundreds of gigabytes. Also, you can edit this option later, and changes will take effect after applying the updated `deploy/cr.yaml` file with `kubectl`.

13.1.5 Enabling compression for backups

There is a possibility to enable [LZ4 compression](#) for backups.

Note

This feature is available only with Percona XtraDB Cluster 8.0 and not Percona XtraDB Cluster 5.7.

To enable compression, use `pxc.configuration` key in the `deploy/cr.yaml` configuration file to supply Percona XtraDB Cluster nodes with two additional `my.cnf` options under its `[sst]` and `[xtrabackup]` sections as follows:

```

pxc:
  image: percona/percona-xtradb-cluster:8.0.19-10.1
  configuration: |
    ...
    [sst]
    xstream-opts=--decompress
    [xtrabackup]
    compress=lz4
    ...

```

When enabled, compression will be used for both backups and [SST](#).

13.1.6 Restore the cluster from a previously saved backup

Backups can be restored not only on the Kubernetes cluster where it was made, but also on any Kubernetes-based environment with the installed Operator.

Backups **cannot be restored** to [emptyDir](#) and [hostPath](#) volumes, but it is possible to make a backup from such storage (i. e., from `emptyDir/hostPath` to S3), and later restore it to a [Persistent Volume](#).

Note

When restoring to a new Kubernetes-based environment, make sure it has a `Secrets` object with the same user passwords as in the original cluster. More details about secrets can be found in [System Users](#).

Following things are needed to restore a previously saved backup:

- Make sure that the cluster is running.
- Find out correct names for the **backup** and the **cluster**. Available backups can be listed with the following command:

```
$ kubectl get pxc-backup
```

Note

Obviously, you can make this check only on the same cluster on which you have previously made the backup.

And the following command will list existing Percona XtraDB Cluster names in the current Kubernetes-based environment:

```
$ kubectl get pxc
```

Restoring without point-in-time recovery

When the correct names for the backup and the cluster are known, backup restoration can be done in the following way.

1. Set appropriate keys in the `deploy/backup/restore.yaml` file.

- set `spec.pxcCluster` key to the name of the target cluster to restore the backup on,
- if you are restoring backup on the *same* Kubernetes-based cluster you have used to save this backup, set `spec.backupName` key to the name of your backup,
- if you are restoring backup on the Kubernetes-based cluster *different* from one you have used to save this backup, set `spec.backupSource` subsection instead of `spec.backupName` field to point on the appropriate PVC, or cloud storage:

PVC volume

The `backupSource` key should contain the storage name (which should be configured in the main CR) and PVC Name:

```
...
backupSource:
  destination: pvc/PVC_VOLUME_NAME
  storageName: pvc
...
```

Note

If you need a [headless Service](#) for the restore Pod (i.e. restoring from a Persistent Volume in a tenant network), mention this in the `metadata.annotations` as follows:

```
...
annotations:
  percona.com/headless-service: "true"
...
```

S3-compatible storage

The `backupSource` key should contain `destination` key equal to the S3 bucket with a special `s3://` prefix, followed by the necessary S3 configuration keys, same as in `deploy/cr.yaml` file:

```
...
backupSource:
  destination: s3://S3-BUCKET-NAME/BACKUP-NAME
  s3:
    credentialsSecret: my-cluster-name-backup-s3
    region: us-west-2
    endpointUrl: https://URL-OF-THE-S3-COMPATIBLE-STORAGE
...
```

Azure Blob storage

The `backupSource` key should contain `destination` key equal to the Azure Blob container and backup name, followed by the necessary Azure configuration keys, same as in `deploy/cr.yaml` file:

```
...
backupSource:
  destination: AZURE-CONTAINER-NAME/BACKUP-NAME
  azure:
    credentialsSecret: my-cluster-azure-secret
...
```

2. After that, the actual restoration process can be started as follows:

```
$ kubectl apply -f deploy/backup/restore.yaml
```

Note

Storing backup settings in a separate file can be replaced by passing its content to the `kubectl apply` command as follows:

```
$ cat <<EOF | kubectl apply -f-
apiVersion: "pxc.percona.com/v1"
kind: "PerconaXtraDBClusterRestore"
metadata:
  name: "restore1"
spec:
  pxcCluster: "cluster1"
```

```
backupName: "backup1"
EOF
```

Restoring backup with point-in-time recovery

Note

Disable the point-in-time functionality on the existing cluster before restoring a backup on it, regardless of whether the backup was made with point-in-time recovery or without it.

If the point-in-time recovery feature was enabled, you can put additional restoration parameters to the `restore.yaml` file `pitr` section for the most fine-grained restoration.

- `backupSource` key should contain `destination` key equal to the s3 bucket with a special `s3://` prefix, followed by the necessary S3 configuration keys, same as in `deploy/cr.yaml` file: `s3://S3-BUCKET-NAME/BACKUP-NAME`,
- `type` key can be equal to one of the following options,
 - `date` - roll back to specific date,
 - `transaction` - roll back to specific transaction (available since Operator 1.8.0),
 - `latest` - recover to the latest possible transaction,
 - `skip` - skip a specific transaction (available since Operator 1.7.0).
- `date` key is used with `type=date` option - it contains value in datetime format,
- `gtid` key (available since Operator 1.8.0) or `gtidSet` key (with Operator 1.7.0) used with `type=transaction` option - it contains exact GTID or GTIDSet (the restore will not include the transaction with specified GTID, but the one before it),
- if you have necessary backup storage mentioned in the `backup.storages` subsection of the `deploy/cr.yaml` configuration file, you can just set `backupSource.storageName` key in the `deploy/backup/restore.yaml` file to the name of the appropriate storage,
- if there is no necessary backup storage in `deploy/cr.yaml`, set your storage details instead of using the `backupSource.storageName` field:

```
...
backupSource:
  s3:
    bucket: S3-BUCKET-NAME
    credentialsSecret: my-cluster-name-backup-s3
    endpointUrl: https://URL-OF-THE-S3-COMPATIBLE-STORAGE
    region: us-west-2
...
```

The resulting `restore.yaml` file may look as follows:

```
apiVersion: pxc.percona.com/v1
kind: PerconaXtraDBClusterRestore
metadata:
  name: restore1
spec:
  pxcCluster: cluster1
  backupName: backup1
  pitr:
    type: date
    date: "2020-12-31 09:37:13"
    backupSource:
      storageName: "s3-us-west"
```

The actual restoration process can be started as follows:

```
$ kubectl apply -f deploy/backup/restore.yaml
```

 **Note**

Storing backup settings in a separate file can be replaced by passing its content to the `kubectl apply` command as follows:

```
$ cat <<EOF | kubectl apply -f-
apiVersion: "pxc.percona.com/v1"
kind: "PerconaXtraDBClusterRestore"
metadata:
  name: "restore1"
spec:
  pxcCluster: "cluster1"
  backupName: "backup1"
  pitr:
    type: date
    date: "2020-12-31 09:37:13"
  backupSource:
    storageName: "s3-us-west"
EOF
```

Take into account, that Operator monitors the binlog gaps detected by binlog collector, if any. If backup contains such gaps, the Operator will mark the status of the latest successful backup with a new condition field that indicates backup can't guarantee consistent point-in-time recovery. This condition looks as follows:

```
apiVersion: pxc.percona.com/v1
kind: PerconaXtraDBClusterBackup
metadata:
  name: backup1
spec:
  pxcCluster: pitr
  storageName: minio
status:
  completed: "2022-11-25T15:57:29Z"
  conditions:
  - lastTransitionTime: "2022-11-25T15:57:48Z"
    message: Binlog with GTID set e41eb219-6cd8-11ed-94c8-9ebf697d3d20:21-22 not found
    reason: BinlogGapDetected
    status: "False"
    type: PITRReady
    state: Succeeded
```

Trying to restore from such backup (with the condition value "False") with point-in-time recovery will result in the following error:

```
Backup doesn't guarantee consistent recovery with PITR. Annotate PerconaXtraDBClusterRestore with percona.com/unsafe-pitr to force it.
```

You can disable this check and force the restore by annotating it with `pxc.percona.com/unsafe-pitr` as follows:

```
apiVersion: pxc.percona.com/v1
kind: PerconaXtraDBClusterRestore
metadata:
  annotations:
    percona.com/unsafe-pitr: "true"
  name: restore2
spec:
  pxcCluster: pitr
  backupName: backup1
  pitr:
    type: latest
  backupSource:
    storageName: "minio-binlogs"
```

13.1.7 Delete the unneeded backup

The maximum amount of stored backups is controlled by the `backup.schedule.keep` option (only successful backups are counted). Older backups are automatically deleted, so that amount of stored backups do not exceed this number. Setting `keep=0` or removing this option from `deploy/cr.yaml` disables automatic deletion of backups.

Manual deleting of a previously saved backup requires not more than the backup name. This name can be taken from the list of available backups returned by the following command:

```
$ kubectl get pxc-backup
```

When the name is known, backup can be deleted as follows:

```
$ kubectl delete pxc-backup/<backup-name>
```

13.1.8 Copy backup to a local machine

Make a local copy of a previously saved backup requires not more than the backup name. This name can be taken from the list of available backups returned by the following command:

```
$ kubectl get pxc-backup
```

When the name is known, backup can be downloaded to the local machine as follows:

```
$ ./deploy/backup/copy-backup.sh <backup-name> path/to/dir
```

For example, this downloaded backup can be restored to the local installation of Percona Server:

```
$ service mysqld stop
$ rm -rf /var/lib/mysql/*
$ cat xtrabackup.stream | xbstream -x -C /var/lib/mysql
$ xtrabackup --prepare --target-dir=/var/lib/mysql
$ chown -R mysql:mysql /var/lib/mysql
$ service mysqld start
```

Last update: 2022-12-07

13.2 Upgrade Database and Operator

Starting from the version 1.1.0, Percona Operator for MySQL based on Percona XtraDB Cluster allows upgrades to newer versions. The upgradable components of the cluster are the following ones: * the Operator; * [Custom Resource Definition \(CRD\)](#), * Database Management System (Percona XtraDB Cluster).

The list of recommended upgrade scenarios includes two variants:

- Upgrade to the new versions of the Operator *and* Percona XtraDB Cluster,
- Minor Percona XtraDB Cluster version upgrade *without* the Operator upgrade.

13.2.1 Upgrade to the new versions of the Operator *and* Percona XtraDB Cluster

In this scenario, components of the cluster are upgraded in the following order:

1. The Operator and CRD,
2. Percona XtraDB Cluster.

Upgrading the Operator and CRD

 **Note**

The Operator supports **last 3 versions of the CRD**, so it is technically possible to skip upgrading the CRD and just upgrade the Operator. If the CRD is older than the new Operator version *by no more than three releases*, you will be able to continue using the old CRD and even carry on Percona XtraDB Cluster minor version upgrades with it. But the recommended way is to update the Operator *and* CRD.

 **Note**

Only the incremental update to a nearest minor version of the Operator is supported (for example, update from 1.4.0 to 1.5.0). To update to a newer version, which differs from the current version by more than one, make several incremental updates sequentially.

MANUAL UPGRADE

The upgrade includes the following steps.

1. Update the [Custom Resource Definition](#) for the Operator, taking it from the official repository on Github, and do the same for the Role-based access control:

```
$ kubectl apply -f https://raw.githubusercontent.com/percona/percona-xtradb-cluster-operator/v1.12.0/deploy/crd.yaml
$ kubectl apply -f https://raw.githubusercontent.com/percona/percona-xtradb-cluster-operator/v1.12.0/deploy/rbac.yaml
```

2. Now you should [apply a patch](#) to your deployment, supplying necessary image name with a newer version tag. You can find the proper image name for the current Operator release [in the list of certified images](#) (for older releases, please refer to the [old releases documentation archive](#)). For example, updating to the 1.12.0 version should look as follows.

```
$ kubectl patch deployment percona-xtradb-cluster-operator \
-p '{"spec":{"template":{"spec":{"containers":[{"name":"percona-xtradb-cluster-operator","image":"percona/percona-xtradb-cluster-operator:1.12.0"}]}}}}'
```

3. The deployment rollout will be automatically triggered by the applied patch. You can track the rollout process in real time with the `kubectl rollout status` command with the name of your cluster:

```
$ kubectl rollout status deployments percona-xtradb-cluster-operator
```

 **Note**

Labels set on the Operator Pod will not be updated during upgrade.

UPGRADE VIA HELM

If you have [installed the Operator using Helm](#), you can upgrade the Operator with the `helm upgrade` command.

1. In case if you installed the Operator with no [customized parameters](#), the upgrade can be done as follows:

```
$ helm upgrade my-op percona/pxc-operator --version 1.12.0
```

The `my-op` parameter in the above example is the name of a [release object](#) which which you have chosen for the Operator when installing its Helm chart.

If the Operator was installed with some [customized parameters](#), you should list these options in the upgrade command.

 **Note**

You can get list of used options in YAML format with the `helm get values my-op -a > my-values.yaml` command, and this file can be directly passed to the upgrade command as follows:

```
$ helm upgrade my-op percona/pxc-operator --version 1.12.0 -f my-values.yaml
```

2. Update the [Custom Resource Definition](#) for the Operator, taking it from the official repository on Github, and do the same for the Role-based access control:

```
$ kubectl apply -f https://raw.githubusercontent.com/percona/percona-xtradb-cluster-operator/v1.12.0/deploy/crd.yaml
$ kubectl apply -f https://raw.githubusercontent.com/percona/percona-xtradb-cluster-operator/v1.12.0/deploy/rbac.yaml
```

 **Note**

You can use `helm upgrade` to upgrade the Operator only. The Database Management System (Percona XtraDB Cluster) should be upgraded in the same way whether you used helm to install it or not.

Upgrading Percona XtraDB Cluster

The recommended way to update the database management system (Percona XtraDB Cluster) is to use the *Smart Update strategy*, which means that upgrade process is controlled by the Operator. But it is also possible to [carry updates manually](#), if needed.

In case of Smart Updates, the Operator can either detect the availability of the Percona XtraDB Cluster version or rely on the user's choice of the version. In both cases Pods are restarted by the Operator automatically in the order, which assures the primary instance to be updated last, preventing possible connection issues until the whole cluster is updated to the new settings.

To have this upgrade method enabled, make sure that the `updateStrategy` key in the `deploy/cr.yaml` configuration file is set to `SmartUpdate`.

When automatic updates are enabled, the Operator will carry on upgrades according to the following algorithm. It will query a special *Version Service* server at scheduled times to obtain fresh information about version numbers and valid image paths needed for the upgrade. If the current version should be upgraded, the Operator updates the CR to reflect the new image paths and carries on sequential Pods deletion in a safe order, allowing StatefulSet to redeploy the cluster Pods with the new image.

The upgrade details are set in the `upgradeOptions` section of the `deploy/cr.yaml` configuration file. Make the following edits to configure updates:

1. Set the `apply` option to one of the following values:

- `Recommended` - automatic upgrades will choose the most recent version of software flagged as Recommended (for clusters created from scratch, the Percona XtraDB Cluster 8.0 version will be selected instead of the Percona XtraDB Cluster 5.7 one regardless of the image path; for already existing clusters, the 8.0 vs. 5.7 branch choice will be preserved),
- `8.0-recommended`, `5.7-recommended` - same as above, but preserves specific major Percona XtraDB Cluster version for newly provisioned clusters (ex. 8.0 will not be automatically used instead of 5.7),
- `Latest` - automatic upgrades will choose the most recent version of the software available,
- `8.0-latest`, `5.7-latest` - same as above, but preserves specific major Percona XtraDB Cluster version for newly provisioned clusters (ex. 8.0 will not be automatically used instead of 5.7),
- `version number` - specify the desired version explicitly (version numbers are specified as `8.0.29-21.1`, `5.7.39-31.61`, etc.),
- `Never` or `Disabled` - disable automatic upgrades (the default choice starting from the Operator 1.12.0)

Note

When automatic upgrades are disabled by the `apply` option, Smart Update functionality will continue working for changes triggered by other events, such as updating a ConfigMap, rotating a password, or changing resource values. So, Smart Update strategy can be used even if the automatic update by querying the Version Service server is disabled.

2. Make sure the `versionServiceEndpoint` key is set to a valid Version Server URL (otherwise Smart Updates will not occur).

- You can use the URL of the official Percona's Version Service (default). Set `versionServiceEndpoint` to `https://check.percona.com`.
- Alternatively, you can run Version Service inside your cluster. This can be done with the `kubectl` command as follows:

```
$ kubectl run version-service --image=perconalab/version-service --env="SERVE_HTTP=true" --port 11000 --expose
```

Note

Version Service is never checked if automatic updates are disabled. If automatic updates are enabled, but Version Service URL can not be reached, upgrades will not occur.

3. Use the `schedule` option to specify the update checks time in CRON format.

The following example sets the midnight update checks with the official Percona's Version Service:

```
spec:
  updateStrategy: SmartUpdate
  upgradeOptions:
    apply: Recommended
    versionServiceEndpoint: https://check.percona.com
    schedule: "0 0 * * *"
  ...
```

13.2.2 Minor Percona XtraDB Cluster version upgrade without the Operator upgrade

In this scenario, the database management system (Percona XtraDB Cluster) is the only component of the cluster which is upgraded, and minor version upgrade is the only one to occur. For example, the image `percona-xtradb-cluster:8.0.25-15.1` can be upgraded to `percona-xtradb-cluster:8.0.27-18.1`.

You can find the proper image name for the current Operator release [in the list of certified images](#). For older releases, please refer to the [old releases documentation archive](#).

Last update: 2022-12-07

13.3 Scale Percona XtraDB Cluster on Kubernetes and OpenShift

One of the great advantages brought by Kubernetes and the OpenShift platform is the ease of an application scaling. Scaling an application results in adding or removing the Pods and scheduling them to available Kubernetes nodes.

Size of the cluster is controlled by a [size key](#) in the [Custom Resource options](#) configuration. That's why scaling the cluster needs nothing more but changing this option and applying the updated configuration file. This may be done in a specifically saved config, or on the fly, using the following command:

```
$ kubectl scale --replicas=5 pxc/cluster1
```

In this example we have changed the size of the Percona XtraDB Cluster to 5 instances.

13.3.1 Increase the Persistent Volume Claim size

Kubernetes manages storage with a PersistentVolume (PV), a segment of storage supplied by the administrator, and a PersistentVolumeClaim (PVC), a request for storage from a user. In Kubernetes v1.11 the feature was added to allow a user to increase the size of an existing PVC object. The user cannot shrink the size of an existing PVC object. Certain volume types support, by default, expanding PVCs (details about PVCs and the supported volume types can be found in [Kubernetes documentation](#))

The following are the steps to increase the size:

1. Extract and backup the yaml file for the cluster

```
$ kubectl get pxc cluster1 -o yaml --export > CR_backup.yaml
```

1. Now you should delete the cluster.

Warning

Make sure that `delete-pxc-pvc` finalizer is not set in your custom resource, **otherwise all cluster data will be lost!**

You can use the following command to delete the cluster:

```
$ kubectl delete -f CR_backup.yaml
```

2. For each node, edit the yaml to resize the PVC object.

```
$ kubectl edit pvc datadir-cluster1-pxc-0
```

In the yaml, edit the `spec.resources.requests.storage` value.

```
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 6Gi
```

Perform the same operation on the other nodes.

```
$ kubectl edit pvc datadir-cluster1-pxc-1
$ kubectl edit pvc datadir-cluster1-pxc-2
```

3. In the CR configuration file, use vim or another text editor to edit the PVC size.

```
$ vim CR_backup.yaml
```


4. Apply the updated configuration to the cluster.

```
$ kubectl apply -f CR_backup.yaml
```

Last update: 2022-12-04

13.4 Set up Percona XtraDB Cluster cross-site replication

The cross-site replication involves configuring one Percona XtraDB Cluster as *Source*, and another Percona XtraDB Cluster as *Replica* to allow an asynchronous replication between them:



The Operator automates configuration of *Source* and *Replica* Percona XtraDB Clusters, but the feature itself is not bound to Kubernetes. Either *Source* or *Replica* can run outside of Kubernetes, be regular MySQL and be out of the Operators' control.

This feature can be useful in several cases: for example, it can simplify migration from on-premises to the cloud with replication, and it can be really helpful in case of the disaster recovery too.

Note

Cross-site replication is based on [Automatic Asynchronous Replication Connection Failover](#). Therefore it requires MySQL 8.0.22+ (Percona XtraDB Cluster 8.0.22+) to work.

Setting up MySQL for asynchronous replication without the Operator is described [here](#) and is out of the scope for this document.

Configuring the cross-site replication for the cluster controlled by the Operator is explained in the following subsections.

13.4.1 Configuring cross-site replication on Source instances

You can configure *Source* instances for cross-site replication with `spec.pxc.replicationChannels` subsection in the `deploy/cr.yaml` configuration file. It is an array of channels, and you should provide the following keys for the channel in your *Source* Percona XtraDB Cluster:

- `pxc.replicationChannels.[].name` key is the name of the channel,
- `pxc.replicationChannels.[].isSource` key should be set to `true`.

Here is an example:

```
spec:
  pxc:
    replicationChannels:
      - name: pxc1_to_pxc2
        isSource: true
```

The cluster will be ready for asynchronous replication when you apply changes as usual:

```
$ kubectl apply -f deploy/cr.yaml
```

13.4.2 Exposing instances of Percona XtraDB Cluster

You need to expose every Percona XtraDB Cluster Pod of the *Source* cluster to make it possible for the *Replica* cluster to connect. This is done through the `pxc.expose` section in the `deploy/cr.yaml` configuration file as follows.

```
spec:
  pxc:
    expose:
      enabled: true
      type: LoadBalancer
```

Note

This will create a LoadBalancer per each Percona XtraDB Cluster Pod. In most cases, for cross-region replication to work this Load Balancer should be internet-facing.

To list the endpoints assigned to PXC Pods list the Kubernetes Service objects by executing `kubectl get services -l "app.kubernetes.io/instance=CLUSTER_NAME"` command.

13.4.3 Configuring cross-site replication on Replica instances

You can configure *Replica* instances for cross-site replication with `spec.pxc.replicationChannels` subsection in the `deploy/cr.yaml` configuration file. It is an array of channels, and you should provide the following keys for the channel in your *Replica* Percona XtraDB Cluster:

- `pxc.replicationChannels.[].name` key is the name of the channel,
- `pxc.replicationChannels.[].isSource` key should be set to `false`,
- `pxc.replicationChannels.[].sourcesList` is the list of *Source* cluster names from which Replica should get the data,
- `pxc.replicationChannels.[].sourcesList.[].host` is the host name or IP address of the Source,
- `pxc.replicationChannels.[].sourcesList.[].port` is the port of the source (`3306` port will be used if nothing specified),
- `pxc.replicationChannels.[].sourcesList.[].weight` is the *weight* of the source (in the event of a connection failure, a new source is selected from the list based on a weighted priority).

Here is the example:

```
spec:
  pxc:
    replicationChannels:
      - name: uspxc1_to_pxc2
        isSource: false
        sourcesList:
          - host: pxc1.source.percona.com
            port: 3306
            weight: 100
          - host: pxc2.source.percona.com
            weight: 100
          - host: pxc3.source.percona.com
            weight: 100
      - name: eu_to_pxc2
        isSource: false
        sourcesList:
          - host: pxc1.source.percona.com
            port: 3306
            weight: 100
          - host: pxc2.source.percona.com
            weight: 100
          - host: pxc3.source.percona.com
            weight: 100
```

The cluster will be ready for asynchronous replication when you apply changes as usual:

```
$ kubectl apply -f deploy/cr.yaml
```

Note

You can also [configure SSL channel for replication](#). Following options allow you using replication over an encrypted channel. Set the `replicationChannels.configuration.ssl` key to true, optionally enable host name identity verification with the `replicationChannels.configuration.sslSkipVerify` key, and set `replicationChannels.configuration.ca` key to the path name of the Certificate Authority (CA) certificate file:

```
replicationChannels:
- isSource: false
  name: uspxc1_to_pxc2
  configuration:
    ssl: true
    sslSkipVerify: true
    ca: '/etc/mysql/ssl/ca.crt'
  ...
```

13.4.4 System user for replication

Replication channel demands a special [system user](#) with same credentials on both *Source* and *Replica*.

The Operator creates a system-level Percona XtraDB Cluster user named `replication` for this purpose, with credentials stored in a Secret object [along with other system users](#).

Note

If the cluster is outside of Kubernetes and is not under the Operator's control, [the appropriate user with necessary permissions](#) should be created manually.

You can change a password for this user as follows:

in Linux

```
$ kubectl patch secret/cluster1-secrets -p '{"data":{"replication": "$(echo -n new_password | base64 --wrap=0)"}'}
```

in macOS

```
$ kubectl patch secret/cluster1-secrets -p '{"data":{"replication": "$(echo -n new_password | base64)"}'}
```

If you have changed the `replication` user's password on the Source cluster, and you use the Operator version 1.9.0, you can have a *replication is not running* error message in log, similar to the following one:

```
{"level":"info","ts":"1629715578.2569592","caller":"zapr/zapr.go:69","msg":"Replication for channel is not running. Please, check the replication status","channel":"pxc2_to_pxc1"}
```

Fixing this involves the following steps.

1. Find the Replica Pod which was chosen by the Operator for replication, using the following command:

```
$ kubectl get pods --selector percona.com/replicationPod=true
```

2. Get the shell access to this Pod and login to the MySQL monitor as a [root user](#):

```
$ kubectl exec -c pxc --stdin --tty <pod_name> -- /bin/bash
bash-4.4$ mysql -uroot -proot_password
```

3. Execute the following three SQL commands to propagate the `replication` user password from the Source cluster to Replica:

```
STOP REPLICATION IO_THREAD FOR CHANNEL '$REPLICATION_CHANNEL_NAME';
CHANGE MASTER TO MASTER_PASSWORD='$NEW_REPLICATION_PASSWORD' FOR CHANNEL '$REPLICATION_CHANNEL_NAME';
START REPLICATION IO_THREAD FOR CHANNEL '$REPLICATION_CHANNEL_NAME';
```

Last update: 2022-12-07

13.5 Monitoring

Percona Monitoring and Management (PMM) [provides an excellent solution](#) to monitor Percona XtraDB Cluster.

Note

Only PMM 2.x versions are supported by the Operator.

PMM is a client/server application. *PMM Client* runs on each node with the database you wish to monitor: it collects needed metrics and sends gathered data to *PMM Server*. As a user, you connect to PMM Server to see database metrics on a number of dashboards.

That's why PMM Server and PMM Client need to be installed separately.

13.5.1 Installing the PMM Server

PMM Server runs as a *Docker image*, a *virtual appliance*, or on an *AWS instance*. Please refer to the [official PMM documentation](#) for the installation instructions.

13.5.2 Installing the PMM Client

The following steps are needed for the PMM client installation in your Kubernetes-based environment:

1. The PMM client installation is initiated by updating the `pmm` section in the `deploy/cr.yaml` file.

- set `pmm.enabled=true`
- set the `pmm.serverHost` key to your PMM Server hostname,
- authorize PMM Client within PMM Server in one of two ways:

with token-based authorization (recommended)

Acquire the API Key from your PMM Server and set `pmmserverkey` in the `deploy/secrets.yaml` secrets file to this obtained API Key value.

with password-based authorization

Check that the `serverUser` key in the `deploy/cr.yaml` file contains your PMM Server user name (`admin` by default), and make sure the `pmmserver` key in the `deploy/secrets.yaml` secrets file contains the password specified for the PMM Server during its installation.

Password-based authorization method is deprecated since the Operator 1.11.0.

Note

You use `deploy/secrets.yaml` file to *create* Secrets Object. The file contains all values for each key/value pair in a convenient plain text format. But the resulting Secrets contain passwords stored as base64-encoded strings. If you want to *update* password field, you'll need to encode the value into base64 format. To do this, you can run `echo -n "password" | base64 --wrap=0` (or just `echo -n "password" | base64` in case of Apple macOS) in your local shell to get valid values. For example, setting the PMM Server API Key to `new_key` in the `cluster1-secrets` object can be done with the following command:

in Linux

```
$ kubectl patch secret/cluster1-secrets -p '{"data":{"pmmserverkey": "'$(echo -n new_key | base64 --wrap=0)'"}}'
```

in macOS

```
$ kubectl patch secret/cluster1-secrets -p '{"data":{"pmmserverkey": "'$(echo -n new_key | base64)'"}}'
```

- you can also use `pmm.pxcParams` and `pmm.proxysqlParams` keys to specify additional parameters for `pmm-admin add mysql` and `pmm-admin add mysql` commands respectively, if needed.

Note

Please take into account that Operator automatically manages common Percona XtraDB Cluster Service Monitoring parameters mentioned in the official PMM documentation, such like username, password, service-name, host, etc. Assigning values to these parameters is not recommended and can negatively affect the functionality of the PMM setup carried out by the Operator.

Apply changes with the `kubectl apply -f deploy/secrets.yaml` command.

When done, apply the edited `deploy/cr.yaml` file:

```
$ kubectl apply -f deploy/cr.yaml
```

2. Check that corresponding Pods are not in a cycle of stopping and restarting. This cycle occurs if there are errors on the previous steps:

```
$ kubectl get pods
$ kubectl logs cluster1-pxc-node-0 -c pmm-client
```

3. Now you can access PMM via `https` in a web browser, with the login/password authentication, and the browser is configured to show Percona XtraDB Cluster metrics.

Last update: 2022-12-04

13.6 Using sidecar containers

The Operator allows you to deploy additional (so-called *sidecar*) containers to the Pod. You can use this feature to run debugging tools, some specific monitoring solutions, etc.

Note

Custom sidecar containers [can easily access other components of your cluster](#).

Therefore they should be used carefully and by experienced users only.

13.6.1 Adding a sidecar container

You can add sidecar containers to Percona XtraDB Cluster, HAProxy, and ProxySQL Pods. Just use `sidecars` subsection in the `pxc`, `haproxy`, or `proxysql` section of the `deploy/cr.yaml` configuration file. In this subsection, you should specify the name and image of your container and possibly a command to run:

```
spec:
  pxc:
    ....
  sidecars:
  - image: busybox
    command: ["/bin/sh"]
    args: ["-c", "while true; do echo echo $(date -u) 'test' >> /dev/null; sleep 5; done"]
    name: my-sidecar-1
    ....
```

Apply your modifications as usual:

```
$ kubectl apply -f deploy/cr.yaml
```

Running `kubectl describe` command for the appropriate Pod can bring you the information about the newly created container:

```
$ kubectl describe pod cluster1-pxc-0
```

Expected output

```
.....
Containers:
  .....
  my-sidecar-1:
    Container ID:   docker://f0c3437295d0ec819753c581aae174a0b8d062337f80897144eb8148249ba742
    Image:          busybox
    Image ID:      docker-pullable://busybox@sha256:139abcf41943b8bcd4bc5c42ee71ddc9402c7ad69ad9e177b0a9bc4541f14924
    Port:          <none>
    Host Port:     <none>
    Command:
    /bin/sh
    Args:
    -c
      while true; do echo echo $(date -u) 'test' >> /dev/null; sleep 5; done
    State:         Running
      Started:     Thu, 11 Nov 2021 10:38:15 +0300
    Ready:         True
    Restart Count: 0
    Environment:   <none>
    Mounts:
    /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-fbrbn (ro)
    .....
```

13.6.2 Getting shell access to a sidecar container

You can login to your sidecar container as follows:

```
$ kubectl exec -it cluster1-pxc-0 -c my-sidecar-1 -- sh
/ #
```

13.6.3 Mount volumes into sidecar containers

It is possible to mount volumes into sidecar containers.

Following subsections describe different [volume types](#), which were tested with sidecar containers and are known to work.

Persistent Volume

You can use [Persistent volumes](#) when you need dynamically provisioned storage which doesn't depend on the Pod lifecycle. To use such volume, you should *claim* durable storage with [persistentVolumeClaim](#) without specifying any non-important details.

The following example requests 1G storage with `sidecar-volume-claim` `PersistentVolumeClaim`, and mounts the correspondent Persistent Volume to the `my-sidecar-1` container's filesystem under the `/volume1` directory:

```
...
sidecars:
- image: busybox
  command: ["/bin/sh"]
  args: ["-c", "while true; do echo echo $(date -u) 'test' >> /dev/null; sleep 5; done"]
  name: my-sidecar-1
  volumeMounts:
  - mountPath: /volume1
    name: sidecar-volume-claim
sidecarPVCs:
- apiVersion: v1
  kind: PersistentVolumeClaim
  metadata:
  name: sidecar-volume-claim
spec:
  resources:
  requests:
  storage: 1Gi
  volumeMode: Filesystem
  accessModes:
  - ReadWriteOnce
```

Secret

You can use a [secret volume](#) to pass the information which needs additional protection (e.g. passwords), to the container. Secrets are stored with the Kubernetes API and mounted to the container as RAM-stored files.

You can mount a secret volume as follows:

```
...
sidecars:
- image: busybox
  command: ["/bin/sh"]
  args: ["-c", "while true; do echo echo $(date -u) 'test' >> /dev/null; sleep 5; done"]
  name: my-sidecar-1
  volumeMounts:
  - mountPath: /secret
    name: sidecar-secret
sidecarVolumes:
- name: sidecar-secret
  secret:
  secretName: mysecret
```

The above example creates a `sidecar-secret` volume (based on already existing `mysecret` [Secret object](#)) and mounts it to the `my-sidecar-1` container's filesystem under the `/secret` directory.

Note

Don't forget you need to [create a Secret Object](#) before you can use it.

configMap

You can use a [configMap volume](#) to pass some configuration data to the container. Secrets are stored with the Kubernetes API and mounted to the container as RAM-stored files.

You can mount a configMap volume as follows:

```
...
sidecars:
- image: busybox
  command: ["/bin/sh"]
  args: ["-c", "while true; do echo echo $(date -u) 'test' >> /dev/null; sleep 5; done"]
  name: my-sidecar-1
  volumeMounts:
  - mountPath: /config
    name: sidecar-config
sidecarVolumes:
- name: sidecar-config
  configMap:
  name: myconfigmap
```

The above example creates a `sidecar-config` volume (based on already existing `myconfigmap` `configMap` object) and mounts it to the `my-sidecar-1` container's filesystem under the `/config` directory.

 **Note**

Don't forget you need to [create a configMap Object](#) before you can use it.

Last update: 2022-12-04

13.7 Pause/resume Percona XtraDB Cluster

There may be external situations when it is needed to shutdown the Percona XtraDB Cluster for a while and then start it back up (some works related to the maintenance of the enterprise infrastructure, etc.).

The `deploy/cr.yaml` file contains a special `spec.pause` key for this. Setting it to `true` gracefully stops the cluster:

```
spec:
  .....
  pause: true
```

Pausing the cluster may take some time, and when the process is over, you will see only the Operator Pod running:

```
$ kubectl get pods
NAME                                READY   STATUS    RESTARTS   AGE
percona-xtradb-cluster-operator-79966668bd-rswbk 1/1     Running   0           12m
```

To start the cluster after it was shut down just revert the `spec.pause` key to `false`.

Starting the cluster will take time. The process is over when all Pods have reached their Running status:

```
NAME                                READY   STATUS    RESTARTS   AGE
cluster1-haproxy-0                 2/2     Running   0           6m17s
cluster1-haproxy-1                 2/2     Running   0           4m59s
cluster1-haproxy-2                 2/2     Running   0           4m36s
cluster1-pxc-0                     3/3     Running   0           6m17s
cluster1-pxc-1                     3/3     Running   0           5m3s
cluster1-pxc-2                     3/3     Running   0           3m56s
percona-xtradb-cluster-operator-79966668bd-rswbk 1/1     Running   0           9m54s
```

Last update: 2022-12-04

13.8 Crash Recovery

13.8.1 What does the full cluster crash mean?

A full cluster crash is a situation when all database instances were shut down in random order. Being rebooted after such situation, Pod is continuously restarting, and generates the following errors in the log:

```
It may not be safe to bootstrap the cluster from this node. It was not the last one to leave the cluster and may not contain all the updates.
To force cluster bootstrap with this node, edit the grastate.dat file manually and set safe_to_bootstrap to 1
```

Note

To avoid this, shutdown your cluster correctly as it is written in [Pause/resume Percona XtraDB Cluster](#).

The Percona Operator for MySQL based on Percona XtraDB Cluster provides two ways of recovery after a full cluster crash.

The Operator is providing automatic crash recovery (by default) and semi-automatic recovery starting from the version 1.7. For the previous Operator versions, crash recovery can be done manually.

13.8.2 Automatic Crash Recovery

Crash recovery can be done automatically. This behavior is controlled by the `pxc.autoRecovery` option in the `deploy/cr.yaml` configuration file.

The default value for this option is `true`, which means that automatic recovery is turned on.

If this option is set to `false`, automatic crash recovery is not done, but semi-automatic recovery is still possible.

In this case you need to get the log from `pxc` container from all Pods using the following command:

```
$ for i in $(seq 0 $(($(kubectl get pxc cluster1 -o jsonpath='{.spec.pxc.size}')-1)); do echo "#####cluster1-pxc-$i#####"; kubectl logs cluster1-pxc-$i -c pxc | grep '(seqno):' ; done
```

The output of this command should be similar to the following one:

```
#####cluster1-pxc-0#####
It is cluster1-pxc-0.cluster1-pxc.default.svc.cluster.local node with sequence number (seqno): 18
#####cluster1-pxc-1#####
It is cluster1-pxc-1.cluster1-pxc.default.svc.cluster.local node with sequence number (seqno): 18
#####cluster1-pxc-2#####
It is cluster1-pxc-2.cluster1-pxc.default.svc.cluster.local node with sequence number (seqno): 19
```

Now find the Pod with the largest `seqno` (it is `cluster1-pxc-2` in the above example).

Now execute the following commands to start this instance:

```
$ kubectl exec cluster1-pxc-2 -c pxc -- sh -c 'kill -s USR1 1'
```

13.8.3 Manual Crash Recovery

Warning

This method includes a lot of operations, and therefore, it is intended for advanced users only!

This method involves the following steps:

- swap the original Percona XtraDB Cluster image with the [debug image](#), which does not reboot after the crash, and force all Pods to run it,
- find the Pod with the most recent Percona XtraDB Cluster data, run recovery on it, start `mysqld`, and allow the cluster to be restarted,
- revert all temporary substitutions.

Let's assume that a full crash did occur for the cluster named `cluster1`, which is based on three Percona XtraDB Cluster Pods.

Note

The following commands are written for Percona XtraDB Cluster 8.0. The same steps are also for Percona XtraDB Cluster 5.7 unless specifically indicated otherwise.

1. Check the current Update Strategy with the following command to make sure [Smart Updates](#) are turned off during the recovery:

```
$ kubectl get pxc cluster1 -o jsonpath='{.spec.updateStrategy}'
```

If the returned value is `SmartUpdate`, please change it to `onDelete` with the following command:

```
$ kubectl patch pxc cluster1 --type=merge --patch '{"spec": {"updateStrategy": "OnDelete" }}'
```

2. Change the normal PXC image inside the cluster object to the debug image:

Note

Please make sure the Percona XtraDB Cluster version for the debug image matches the version currently in use in the cluster. You can run the following command to find out which Percona XtraDB Cluster image is in use:

```
$ kubectl get pxc cluster1 -o jsonpath='{.spec.pxc.image}'
```

```
$ kubectl patch pxc cluster1 --type="merge" -p '{"spec":{"pxc":{"image":"percona/percona-xtradb-cluster:8.0.29-21.1-debug"}}}'
```

Note

For Percona XtraDB Cluster 5.7 this command should be as follows:

```
$ kubectl patch pxc cluster1 --type="merge" -p '{"spec":{"pxc":{"image":"percona/percona-xtradb-cluster:5.7.39-31.61-debug"}}}'
```

1. Restart all Pods:

```
$ for i in $(seq 0 $(($(kubectl get pxc cluster1 -o jsonpath='{.spec.pxc.size}'-1))); do kubectl delete pod cluster1-pxc-$i --force --grace-period=0; done
```

2. Wait until the Pod 0 is ready, and execute the following code (it is required for the Pod liveness check):

```
$ for i in $(seq 0 $(($(kubectl get pxc cluster1 -o jsonpath='{.spec.pxc.size}'-1))); do until [[ $(kubectl get pod cluster1-pxc-$i -o jsonpath='{.status.phase}') == 'Running' ]]; do sleep 10; done; kubectl exec cluster1-pxc-$i -- touch /var/lib/mysql/sst_in_progress; done
```

3. Wait for all Percona XtraDB Cluster Pods to start, and execute the following code to make sure no mysqld processes are running:

```
$ for i in $(seq $(($(kubectl get pxc cluster1 -o jsonpath='{.spec.pxc.size}'-1))); do pid=$(kubectl exec cluster1-pxc-$i -- ps -C mysqld-ps -o pid=); if [[ -n "$pid" ]]; then kubectl exec cluster1-pxc-$i -- kill -9 $pid; fi; done
```

4. Wait for all Percona XtraDB Cluster Pods to start, then find the Percona XtraDB Cluster instance with the most recent data - i.e. the one with the highest [sequence number \(seqno\)](#):

```
$ for i in $(seq 0 $(($(kubectl get pxc cluster1 -o jsonpath='{.spec.pxc.size}'-1))); do echo "#####cluster1-pxc-$i#####"; kubectl exec cluster1-pxc-$i -- cat /var/lib/mysql/grastate.dat; done
```

The output of this command should be similar to the following one:

```
#####cluster1-pxc-0#####
# GALERA saved state
version: 2.1
uid: 7e037079-6517-11ea-a558-8e77af893c93
seqno: 18
safe_to_bootstrap: 0
#####cluster1-pxc-1#####
# GALERA saved state
version: 2.1
uid: 7e037079-6517-11ea-a558-8e77af893c93
seqno: 18
safe_to_bootstrap: 0
#####cluster1-pxc-2#####
# GALERA saved state
version: 2.1
uid: 7e037079-6517-11ea-a558-8e77af893c93
seqno: 19
safe_to_bootstrap: 0
```

Now find the Pod with the largest `seqno` (it is `cluster1-pxc-2` in the above example).

5. Now execute the following commands *in a separate shell* to start this instance:

```
$ kubectl exec cluster1-pxc-2 -- mysqld --wsrep_recover
$ kubectl exec cluster1-pxc-2 -- sed -i 's/safe_to_bootstrap: 0/safe_to_bootstrap: 1/g' /var/lib/mysql/grastate.dat
$ kubectl exec cluster1-pxc-2 -- sed -i 's/wsrep_cluster_address=.*wsrep_cluster_address=gcomm://\//g' /etc/mysql/node.cnf
$ kubectl exec cluster1-pxc-2 -- mysqld
```

The `mysqld` process will initialize the database once again, and it will be available for the incoming connections.

6. Go back to the previous shell and return the original Percona XtraDB Cluster image because the debug image is no longer needed:

Note

Please make sure the Percona XtraDB Cluster version for the debug image matches the version currently in use in the cluster.

```
$ kubectl patch pxc cluster1 --type="merge" -p '{"spec":{"pxc":{"image":"percona/percona-xtradb-cluster:8.0.29-21.1"}}}'
```

Note

For Percona XtraDB Cluster 5.7 this command should be as follows:

```
$ kubectl patch pxc cluster1 --type=merge -p '{"spec":{"pxc":{"image":"percona/percona-xtradb-cluster:5.7.39-31.61"}}}'
```

1. Restart all Pods besides the `cluster1-pxc-2` Pod (the recovery donor).

```
$ for i in $(seq 0 $(($(kubectl get pxc cluster1 -o jsonpath='{.spec.pxc.size}')-1)); do until [[ $(kubectl get pod cluster1-pxc-$i -o jsonpath='{.status.phase}') = 'Running' ]]; do sleep 10; done; kubectl exec cluster1-pxc-$i -- rm /var/lib/mysql/sst_in_progress; done
$ kubectl delete pods --force --grace-period=0 cluster1-pxc-0 cluster1-pxc-1
```

2. Wait for the successful startup of the Pods which were deleted during the previous step, and finally remove the `cluster1-pxc-2` Pod:

```
$ kubectl delete pods --force --grace-period=0 cluster1-pxc-2
```

3. After the Pod startup, the cluster is fully recovered.

Note

If you have changed the update strategy on the 1st step, don't forget to revert it back to `SmartUpdate` with the following command:

```
$ kubectl patch pxc cluster1 --type=merge --patch '{"spec": {"updateStrategy": "SmartUpdate" }}'
```

Last update: 2022-12-04

13.9 Debug

13.9.1 Cluster-level logging

Cluster-level logging involves collecting logs from all Percona XtraDB Cluster Pods in the cluster to some persistent storage. This feature gives the logs a lifecycle independent of nodes, Pods and containers in which they were collected. Particularly, it ensures that Pod logs from previous failures are available for later review.

Log collector is turned on by the `logcollector.enabled` key in the `deploy/cr.yaml` configuration file (`true` by default).

The Operator collects logs using [Fluent Bit Log Processor](#), which supports many output plugins and has broad forwarding capabilities. If necessary, Fluent Bit filtering and advanced features can be configured via the `logcollector.configuration` key in the `deploy/cr.yaml` configuration file.

Logs are stored for 7 days and then rotated.

Collected logs can be examined using the following command:

```
$ kubectl logs cluster1-pxc-1 -c logs
```

Note

Technically, logs are stored on the same Persistent Volume, which is used with the corresponding Percona XtraDB Cluster Pod. Therefore collected logs can be found in `DATADIR` (`/var/lib/mysql/`).

Note

You can parse output of the logs with [jq JSON processor](#) as follows: `kubectl logs cluster1-pxc-1 -c logs -f | jq -R 'fromjson?'`.

13.9.2 Avoid the restart-on-fail loop for Percona XtraDB Cluster containers

The restart-on-fail loop takes place when the container entry point fails (e.g. `mysqld` crashes). In such a situation, Pod is continuously restarting. Continuous restarts prevent to get console access to the container, and so a special approach is needed to make fixes.

You can prevent such infinite boot loop by putting the Percona XtraDB Cluster containers into the infinity loop *without* starting `mysqld`. This behavior of the container entry point is triggered by the presence of the `/var/lib/mysql/sleep-forever` file.

For example, you can do it for the `pxc` container of an appropriate Percona XtraDB Cluster instance as follows:

```
$ kubectl exec -it cluster1-pxc-0 -c pxc -- sh -c 'touch /var/lib/mysql/sleep-forever'
```

If `pxc` container can't start, you can use `logs` container instead:

```
$ kubectl exec -it cluster1-pxc-0 -c logs -- sh -c 'touch /var/lib/mysql/sleep-forever'
```

The instance will restart automatically and run in its usual way as soon as you remove this file (you can do it with a command similar to the one you have used to create the file, just substitute `touch` to `rm` in it).

13.9.3 Special debug images

For the cases when Pods are failing for some reason or just show abnormal behavior, the Operator can be used with a special *debug images*. Percona XtraDB Cluster debug image has the following specifics:

- it avoids restarting on fail,
- it contains additional tools useful for debugging (sudo, telnet, gdb, etc.),
- it has debug mode enabled for the logs.

There are debug versions for all [Percona XtraDB Cluster images](#): they have same names as normal images with a special `-debug` suffix in their version tag: for example, `percona-xtradb-cluster:8.0.29-21.1-debug`.

To use the debug image instead of the normal one, find the needed image name [in the list of certified images](#) and set it for the proper key in the `deploy/cr.yaml` configuration file. For example, set the following value of the `pxc.image` key to use the Percona XtraDB Cluster debug image:

- `percona/percona-xtradb-cluster:8.0.29-21.1-debug` for Percona XtraDB Cluster 8.0,
- `percona/percona-xtradb-cluster:5.7.39-31.61-debug` for Percona XtraDB Cluster 5.7.

The Pod should be restarted to get the new image.

Note

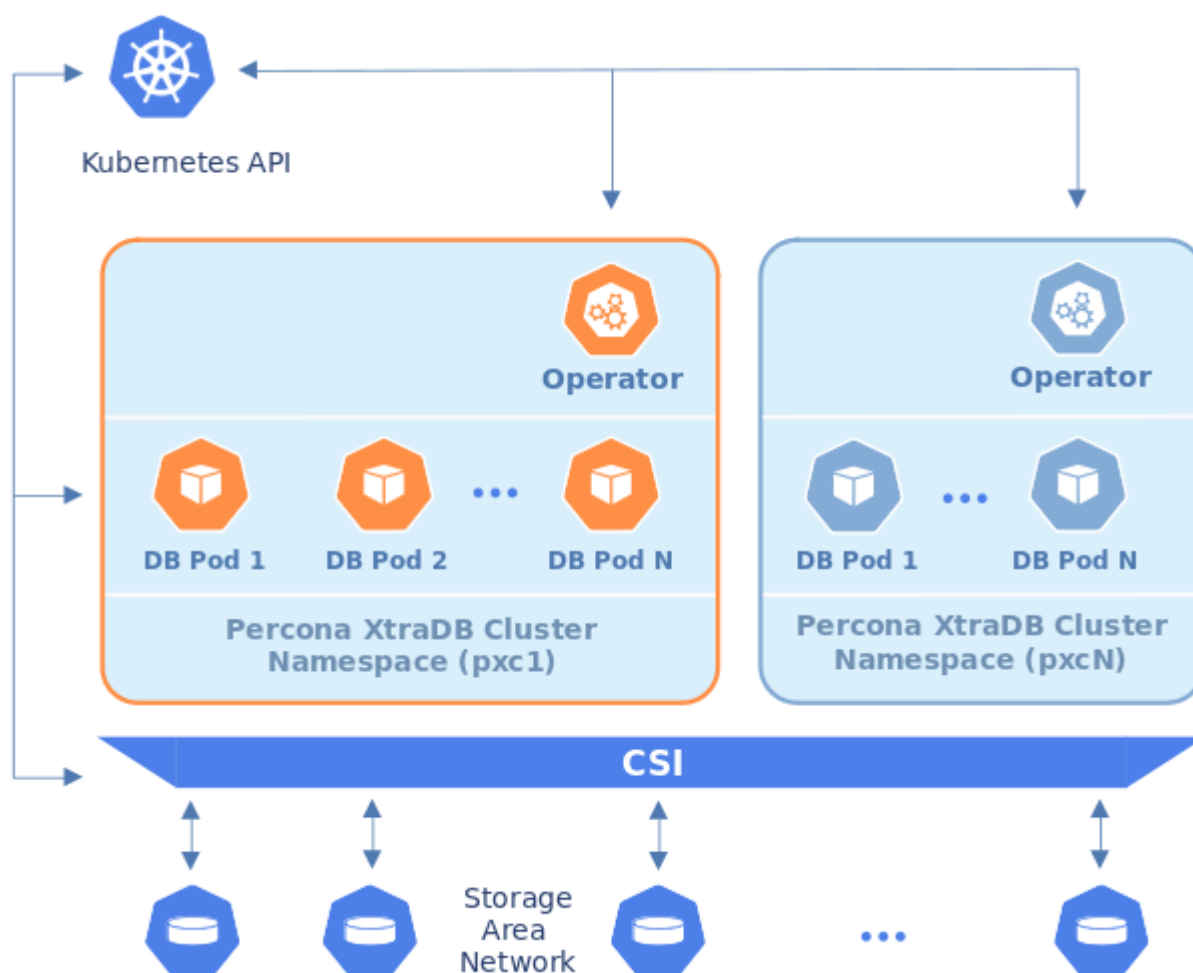
When the Pod is continuously restarting, you may have to delete it to apply image changes.

Last update: 2022-12-04

14. HOWTOs

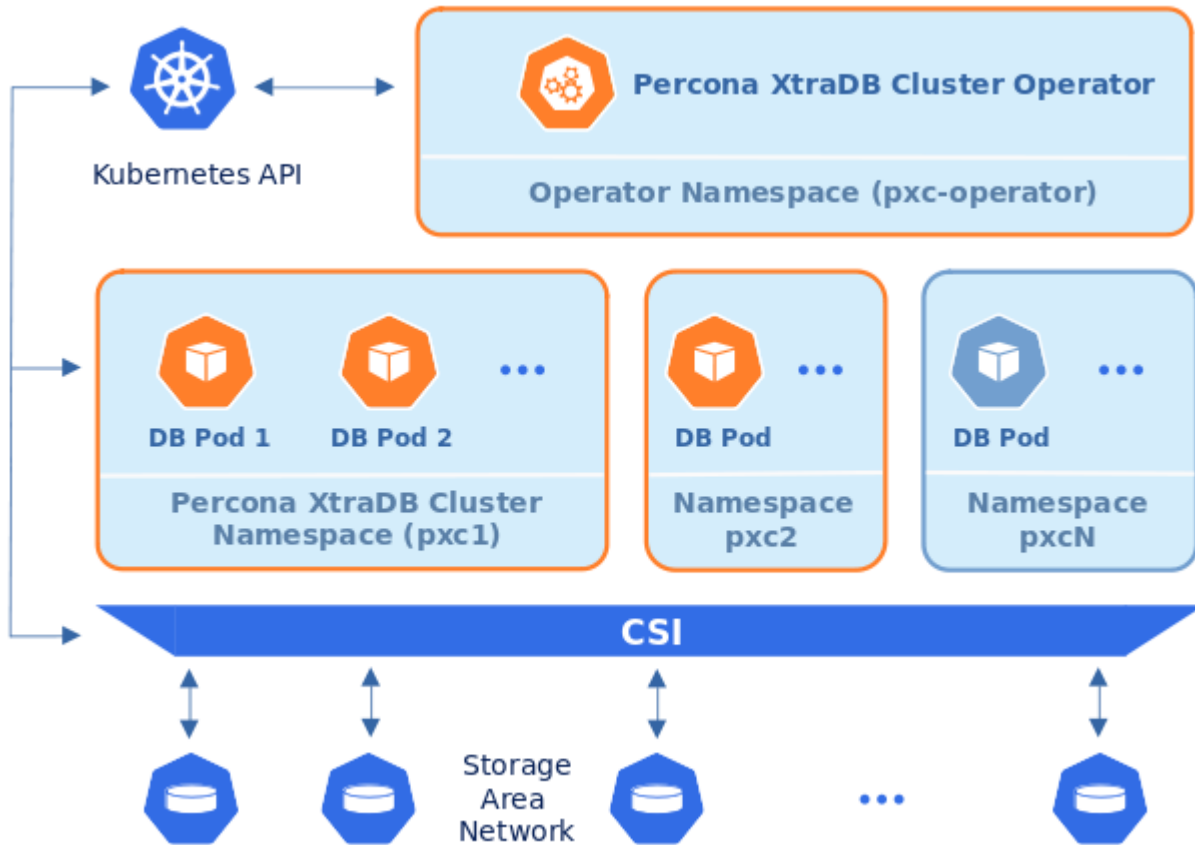
14.1 Install Percona XtraDB Cluster in multi-namespace (cluster-wide) mode

By default, Percona Operator for MySQL based on Percona XtraDB Cluster functions in a specific Kubernetes namespace. You can create one during installation (like it is shown in the [installation instructions](#)) or just use the `default` namespace. This approach allows several Operators to co-exist in one Kubernetes-based environment, being separated in different namespaces:



Still, sometimes it is more convenient to have one Operator watching for Percona XtraDB Cluster custom resources in several namespaces.

We recommend running Percona Operator for MySQL in a traditional way, limited to a specific namespace. But it is possible to run it in so-called *cluster-wide* mode, one Operator watching several namespaces, if needed:



Note

Please take into account that if several Operators are configured to watch the same namespace, it is entirely unpredictable which one will get ownership of the Custom Resource in it, so this situation should be avoided.

To use the Operator in such *cluster-wide* mode, you should install it with a different set of configuration YAML files, which are available in the `deploy` folder and have filenames with a special `cw-` prefix: e.g. `deploy/cw-bundle.yaml`.

While using this cluster-wide versions of configuration files, you should set the following information there:

- `subjects.namespace` option should contain the namespace which will host the Operator,
- `WATCH_NAMESPACE` key-value pair in the `env` section should have `value` equal to a comma-separated list of the namespaces to be watched by the Operator (or just a blank string to make the Operator deal with *all namespaces* in a Kubernetes cluster).

Note

The list of namespaces to watch is fully supported by the Operator starting from the version 1.7 (in the version 1.6 you can only use cluster-wide mode with empty `WATCH_NAMESPACE` key to watch all namespaces). Also, prior to the version 1.12.0 it was necessary to mention the Operator's own namespace in the list of watched namespaces.

The following simple example shows how to install Operator cluster-wide on Kubernetes.

- First of all, clone the `percona-xtradb-cluster-operator` repository:

```
$ git clone -b v1.12.0 https://github.com/percona/percona-xtradb-cluster-operator
$ cd percona-xtradb-cluster-operator
```

- Let's suppose that Operator's namespace should be the `pxc-operator` one. Create it as follows:

```
$ kubectl create namespace pxc-operator
```

Namespaces to be watched by the Operator should be created in the same way if not exist. Let's say the Operator should watch the `pxc` namespace:

```
$ kubectl create namespace pxc
```

- Apply the `deploy/cw-bundle.yaml` file with the following command:

```
$ kubectl apply -f deploy/cw-bundle.yaml -n pxc-operator
```

- After the Operator is started, Percona XtraDB Cluster can be created at any time by applying the `deploy/cr.yaml` configuration file, like in the case of normal installation:

```
$ kubectl apply -f deploy/cr.yaml -n pxc
```

The creation process will take some time. The process is over when both operator and replica set Pods have reached their Running status:

NAME	READY	STATUS	RESTARTS	AGE
cluster1-haproxy-0	2/2	Running	0	6m17s
cluster1-haproxy-1	2/2	Running	0	4m59s
cluster1-haproxy-2	2/2	Running	0	4m36s
cluster1-pxc-0	3/3	Running	0	6m17s
cluster1-pxc-1	3/3	Running	0	5m3s
cluster1-pxc-2	3/3	Running	0	3m56s
percona-xtradb-cluster-operator-7996668bd-rswbk	1/1	Running	0	9m54s

- Check connectivity to newly created cluster

```
$ kubectl run -i --rm --tty percona-client --image=percona:5.7 --restart=Never --env="POD_NAMESPACE=pxc" -- bash -il
percona-client:/$ mysql -h cluster1-proxysql -uroot -proot_password
```

 **Note**

Some Kubernetes-based environments are specifically configured to have communication across Namespaces is not allowed by default network policies. In this case, you should specifically allow the Operator communication across the needed Namespaces. Following the above example, you would need to allow ingress traffic for the `pxc-operator` Namespace from the `pxc` Namespace, and also from the `default` Namespace. You can do it with the NetworkPolicy resource, specified in the YAML file as follows:

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: percona
  namespace: pxc-operator
spec:
  ingress:
    - from:
      - namespaceSelector:
          matchLabels:
            kubernetes.io/metadata.name: pxc
      - namespaceSelector:
          matchLabels:
            kubernetes.io/metadata.name: default
    podSelector: {}
  policyTypes:
    - Ingress
```

Don't forget to apply the resulting file with the usual `kubectl apply` command.

You can find more details about Network Policies [in the official Kubernetes documentation](#).

Last update: 2022-12-07

14.2 How to upgrade Percona XtraDB Cluster manually

Percona Operator for MySQL based on Percona XtraDB Cluster supports upgrades of the database management system (Percona XtraDB Cluster) starting from the Operator version 1.1.0. The Operator 1.5.0 had automated such upgrades with a new upgrade strategy called [Smart Update](#). Smart Update automates the upgrade process and still provides the user full control over updates, so it is the most convenient upgrade strategy.

Still there may be use cases when automatic upgrade of Percona XtraDB Cluster is not an option (for example, you may be using Percona XtraDB Cluster with the Operator version 1.5.0 or earlier), and you have to carry on upgrades manually.

Percona XtraDB Cluster can be upgraded manually using one of the following *upgrade strategies*:

- *Rolling Update*, initiated manually and [controlled by Kubernetes](#),
- *On Delete*, [done by Kubernetes on per-Pod basis](#) when Pods are deleted.

!!! warning:

```
In case of [Smart Updates](update.md#automatic-upgrade), the Operator can either detect the availability of the Percona XtraDB Cluster version or rely on the user's choice of the version. In both cases Pods are restarted by the Operator automatically in the order, which assures the primary instance to be updated last, preventing possible connection issues until the whole cluster is updated to the new settings. Kubernetes-controlled Rolling Update can't guarantee that Pods update order is optimal from the Percona XtraDB Cluster point of view.
```

ROLLING UPDATE STRATEGY AND SEMI-AUTOMATIC UPDATES

Semi-automatic update of Percona XtraDB Cluster can be done as follows:

1. Edit the `deploy/cr.yaml` file, setting `updateStrategy` key to `RollingUpdate`.
2. Now you should [apply a patch](#) to your Custom Resource, setting necessary image names with a newer version tag.

!!! note:

```
Check the version of the Operator you have in your Kubernetes environment. Please refer to the [Operator upgrade guide](update.md#upgrading-the-operator) to upgrade the Operator and CRD, if needed.
```

Patching Custom Resource is done with the `kubectl patch pxc` command. Actual image names can be found [in the list of certified images](#) (for older releases, please refer to the [old releases documentation archive](#)). For example, updating to the 1.12.0 version should look as follows:

For Percona XtraDB Cluster 8.0

```
$ kubectl patch pxc cluster1 --type=merge --patch '{
  "spec": {
    "pxc": { "image": "percona/percona-xtradb-cluster:8.0.29-21.1" },
    "proxysql": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-proxysql" },
    "haproxy": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-haproxy" },
    "backup": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-pxc8.0-backup" },
    "logcollector": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-logcollector" },
    "pmm": { "image": "percona/pmm-client:2.32.0" }
  }
}'
```

For Percona XtraDB Cluster 5.7

```
$ kubectl patch pxc cluster1 --type=merge --patch '{
  "spec": {
    "pxc": { "image": "percona/percona-xtradb-cluster:5.7.39-31.61" },
    "proxysql": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-proxysql" },
    "haproxy": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-haproxy" },
    "backup": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-pxc5.7-backup" },
    "logcollector": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-logcollector" },
    "pmm": { "image": "percona/pmm-client:2.32.0" }
  }
}'
```

3. The deployment rollout will be automatically triggered by the applied patch. You can track the rollout process in real time with the `kubectl rollout status` command with the name of your cluster:

```
$ kubectl rollout status sts cluster1-pxc
```

Manual upgrade (the On Delete strategy)

Manual update of Percona XtraDB Cluster can be done as follows:

1. Edit the `deploy/cr.yaml` file, setting `updateStrategy` key to `OnDelete`.
2. Now you should [apply a patch](#) to your Custom Resource, setting necessary image names with a newer version tag.

!!! note:

Check the version of the Operator you have in your Kubernetes environment. Please refer to the [\[Operator upgrade guide\]\(update.md#upgrading-the-operator\)](#) to upgrade the Operator and CRD, if needed.

Patching Custom Resource is done with the `kubectl patch pxc` command. Actual image names can be found [in the list of certified images](#) (for older releases, please refer to the [old releases documentation archive](#)). For example, updating to the 1.12.0 version should look as follows, depending on whether you are using Percona XtraDB Cluster 5.7 or 8.0.

For Percona XtraDB Cluster 8.0

```
$ kubectl patch pxc cluster1 --type=merge --patch '{
  "spec": {
    "pxc": { "image": "percona/percona-xtradb-cluster:8.0.29-21.1" },
    "proxysql": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-proxysql" },
    "haproxy": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-haproxy" },
    "backup": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-pxc8.0-backup" },
    "logcollector": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-logcollector" },
    "pmm": { "image": "percona/pmm-client:2.32.0" }
  }
}'
```

For Percona XtraDB Cluster 5.7

```
$ kubectl patch pxc cluster1 --type=merge --patch '{
  "spec": {
    "pxc": { "image": "percona/percona-xtradb-cluster:5.7.39-31.61" },
    "proxysql": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-proxysql" },
    "haproxy": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-haproxy" },
    "backup": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-pxc5.7-backup" },
    "logcollector": { "image": "percona/percona-xtradb-cluster-operator:1.12.0-logcollector" },
    "pmm": { "image": "percona/pmm-client:2.32.0" }
  }
}'
```

3. The Pod with the newer Percona XtraDB Cluster image will start after you delete it. Delete targeted Pods manually one by one to make them restart in desired order:
 - a. Delete the Pod using its name with the command like the following one:

```
$ kubectl delete pod cluster1-pxc-2
```

- b. Wait until Pod becomes ready:

```
$ kubectl get pod cluster1-pxc-2
```

The output should be like this:

NAME	READY	STATUS	RESTARTS	AGE
cluster1-pxc-2	1/1	Running	0	3m33s

4. The update process is successfully finished when all Pods have been restarted.

Last update: 2022-12-07

15. Reference

15.1 Custom Resource options

Percona XtraDB Cluster managed by the Operator configured via the spec section of the [deploy/cr.yaml](#) file.

The metadata part of this file contains the following keys:

- `name` (`cluster1` by default) sets the name of your Percona XtraDB Cluster; it should include only [URL-compatible characters](#), not exceed 22 characters, start with an alphabetic character, and end with an alphanumeric character;
- `finalizers.delete-pods-in-order` if present, activates the [Finalizer](#) which controls the proper Pods deletion order in case of the cluster deletion event (on by default).
- `finalizers.delete-pxc-pvc` if present, activates the [Finalizer](#) which deletes [Persistent Volume Claims](#) for Percona XtraDB Cluster Pods after the cluster deletion event (off by default).
- `finalizers.delete-proxysql-pvc` if present, activates the [Finalizer](#) which deletes [Persistent Volume Claim](#) for ProxySQL Pod after the cluster deletion event (off by default).
- `finalizers.delete-ssl` if present, activates the [Finalizer](#) which deletes [objects, created for SSL](#) (Secret, certificate, and issuer) after the cluster deletion event (off by default).

The spec part of the [deploy/cr.yaml](#) file contains the following sections:

Key	Value type	Default	Description
upgradeOptions	subdoc		Percona XtraDB Cluster upgrade options section
pxc	subdoc		Percona XtraDB Cluster general section
haproxy	subdoc		HAProxy section
proxysql	subdoc		ProxySQL section
pmm	subdoc		Percona Monitoring and Management section
backup	subdoc		Percona XtraDB Cluster backups section
allowUnsafeConfigurations	boolean	false	Prevents users from configuring a cluster with unsafe parameters such as starting the cluster with the number of Percona XtraDB Cluster instances which is less than 3, more than 5, or is an even number, with less than 2 ProxySQL or HAProxy Pods, or without TLS/SSL certificates (if false, unsafe parameters will be automatically changed to safe defaults)
enableCRValidationWebhook	boolean	true	Enables or disables schema validation before applying <code>cr.yaml</code> file (works only in cluster-wide mode due to access restrictions)
pause	boolean	false	Pause/resume: setting it to <code>true</code> gracefully stops the cluster, and setting it to <code>false</code> after shut down starts the cluster back
secretsName	string	cluster1-secrets	A name for users secrets
crVersion	string	1.12.0	Version of the Operator the Custom Resource belongs to
ignoreAnnotations	subdoc	iam.amazonaws.com/ role	The list of annotations to be ignored by the Operator
ignoreLabels	subdoc	rack	The list of labels to be ignored by the Operator
vaultSecretName	string	keyring-secret-vault	A secret for the HashiCorp Vault to carry on Data at Rest Encryption
sslSecretName	string	cluster1-ssl	A secret with TLS certificate generated for <i>external</i> communications, see Transport Layer Security (TLS) for details
sslInternalSecretName	string	cluster1-ssl- internal	A secret with TLS certificate generated for <i>internal</i> communications, see Transport Layer Security (TLS) for details
logCollectorSecretName	string	my-log-collector- secrets	A secret for the Fluent Bit Log Collector
initImage	string	percona/percona- xtradb-cluster- operator:1.12.0	An alternative image for the initial Operator installation
tls	subdoc		Extended cert-manager configuration section
updateStrategy	string	SmartUpdate	A strategy the Operator uses for upgrades

15.1.1 Extended cert-manager configuration section

The `tls` section in the `deploy/cr.yaml` file contains various configuration options for additional customization of the [TLS cert-manager](#).

Key	tls.SANs
Value	subdoc
Example	
Description	Additional domains (SAN) to be added to the TLS certificate within the extended cert-manager configuration
Key	tls.issuerConf.name
Value	string
Example	special-selfsigned-issuer
Description	A cert-manager issuer name
Key	tls.issuerConf.kind
Value	string
Example	ClusterIssuer
Description	A cert-manager issuer type
Key	tls.issuerConf.group
Value	string
Example	cert-manager.io
Description	A cert-manager issuer group . Should be <code>cert-manager.io</code> for built-in cert-manager certificate issuers

15.1.2 Upgrade options section

The `upgradeOptions` section in the `deploy/cr.yaml` file contains various configuration options to control Percona XtraDB Cluster upgrades.

Key	<code>upgradeOptions.versionServiceEndpoint</code>
Value	string
Example	<code>https://check.percona.com</code>
Description	The Version Service URL used to check versions compatibility for upgrade
Key	<code>upgradeOptions.apply</code>
Value	string
Example	Disabled
Description	Specifies how updates are processed by the Operator. <code>Never</code> or <code>Disabled</code> will completely disable automatic upgrades, otherwise it can be set to <code>Latest</code> or <code>Recommended</code> or to a specific version string of Percona XtraDB Cluster (e.g. <code>8.0.19-10.1</code>) that is wished to be version-locked (so that the user can control the version running, but use automatic upgrades to move between them)
Key	<code>upgradeOptions.schedule</code>
Value	string
Example	<code>0 2 * * *</code>
Description	Scheduled time to check for updates, specified in the crontab format

15.1.3 PXC section

The `pxc` section in the [deploy/cr.yaml](#) file contains general configuration options for the Percona XtraDB Cluster.

Key	<code>pxc.size</code>
Value	int
Example	3
Description	The size of the Percona XtraDB cluster must be 3 or 5 for High Availability . other values are allowed if the <code>spec.allowUnsafeConfigurations</code> key is set to true
Key	<code>pxc.image</code>
Value	string
Example	<code>percona/percona-xtradb-cluster:8.0.29-21.1</code>
Description	The Docker image of the Percona cluster used (actual image names for Percona XtraDB Cluster 8.0 and Percona XtraDB Cluster 5.7 can be found in the list of certified images)
Key	<code>pxc.autoRecovery</code>
Value	boolean
Example	<code>true</code>
Description	Turns Automatic Crash Recovery on or off
Key	<code>pxc.expose.enabled</code>
Value	boolean
Example	<code>true</code>
Description	Enable or disable exposing Percona XtraDB Cluster nodes with dedicated IP addresses
Key	<code>pxc.expose.type</code>
Value	string
Example	<code>LoadBalancer</code>
Description	The Kubernetes Service Type used for xposure
Key	<code>pxc.expose.trafficPolicy</code>
Value	string
Example	<code>Local</code>
Description	Specifies whether Service should route external traffic to cluster-wide or node-local endpoints (it can influence the load balancing effectiveness)
Key	<code>pxc.expose.loadBalancerSourceRanges</code>
Value	string
Example	<code>10.0.0.0/8</code>
Description	The range of client IP addresses from which the load balancer should be reachable (if not set, there is no limitations)
Key	<code>pxc.expose.annotations</code>
Value	string
Example	<code>networking.gke.io/load-balancer-type: "Internal"</code>

Description	The Kubernetes annotations
Key	pxc.replicationChannels.name
Value	string
Example	<code>pxc1_to_pxc2</code>
Description	Name of the replication channel for cross-site replication
Key	pxc.replicationChannels.isSource
Value	boolean
Example	<code>false</code>
Description	Should the cluster act as Source (<code>true</code>) or Replica (<code>false</code>) in cross-site replication
Key	pxc.replicationChannels.configuration.sourceRetryCount
Value	int
Example	<code>3</code>
Description	Number of retries Replica should do when the existing connection source fails
Key	pxc.replicationChannels.configuration.sourceConnectRetry
Value	int
Example	<code>60</code>
Description	The interval between reconnection attempts in seconds to be used by Replica when the the existing connection source fails
Key	pxc.replicationChannels.configuration.ssl
Value	boolean
Example	<code>false</code>
Description	Turns SSL for replication channels on or off
Key	pxc.replicationChannels.configuration.sslSkipVerify
Value	boolean
Example	<code>true</code>
Description	Turns the host name identity verification for SSL-based replication on or off
Key	pxc.replicationChannels.configuration.ca
Value	string
Example	<code>/etc/mysql/ssl/ca.crt</code>
Description	The path name of the Certificate Authority (CA) certificate file to be used if the SSL for replication channels is turned on
Key	pxc.replicationChannels.sourcesList.host

Value	string
Example	10.95.251.101
Description	For the cross-site replication Replica cluster, this key should contain the hostname or IP address of the Source cluster
Key	pxc.replicationChannels.sourcesList.port
Value	int
Example	3306
Description	For the cross-site replication Replica cluster, this key should contain the Source port number
Key	pxc.replicationChannels.sourcesList.weight
Value	int
Example	100
Description	For the cross-site replication Replica cluster, this key should contain the Source cluster weight (varies from 1 to 100, the cluster with the higher number will be selected as the replication source first)
Key	pxc.readinessDelaySec
Value	int
Example	15
Description	Adds a delay before a run check to verify the application is ready to process traffic
Key	pxc.livenessDelaySec
Value	int
Example	300
Description	Adds a delay before the run check ensures the application is healthy and capable of processing requests
Key	pxc.configuration
Value	string
Example	<pre>[mysqld] wsrep_debug=ON wsrep-provider_options=gcache.size=1G;gcache.recover=yes</pre>
Description	The <code>my.cnf</code> file options to be passed to Percona XtraDB cluster nodes
Key	pxc.imagePullSecrets.name
Value	string
Example	private-registry-credentials
Description	The Kubernetes ImagePullSecret
Key	pxc.priorityClassName
Value	string
Example	high-priority
Description	The Kubernetes Pod priority class

Key	<code>pxc.schedulerName</code>
Value	string
Example	<code>mycustom-scheduler</code>
Description	The Kubernetes Scheduler
Key	<code>pxc.annotations</code>
Value	label
Example	<code>iam.amazonaws.com/role: role-arn</code>
Description	The Kubernetes annotations
Key	<code>pxc.labels</code>
Value	label
Example	<code>rack: rack-22</code>
Description	Labels are key-value pairs attached to objects
Key	<code>pxc.readinessProbes.initialDelaySeconds</code>
Value	int
Example	15
Description	Number of seconds to wait before performing the first readiness probe
Key	<code>pxc.readinessProbes.timeoutSeconds</code>
Value	int
Example	15
Description	Number of seconds after the container has started before readiness probes are initiated
Key	<code>pxc.readinessProbes.periodSeconds</code>
Value	int
Example	30
Description	How often (in seconds) to perform the readiness probe
Key	<code>pxc.readinessProbes.successThreshold</code>
Value	int
Example	1
Description	Minimum consecutive successes for the readiness probe to be considered successful after having failed
Key	<code>pxc.readinessProbes.failureThreshold</code>
Value	int
Example	5
Description	When the readiness probe fails, Kubernetes will try this number of times before marking the Pod Unready
Key	<code>pxc.livenessProbes.initialDelaySeconds</code>

Value	int
Example	300
Description	Number of seconds to wait before performing the first liveness probe
Key	pxc.livenessProbes.timeoutSeconds
Value	int
Example	5
Description	Number of seconds after the container has started before liveness probes are initiated
Key	pxc.livenessProbes.periodSeconds
Value	int
Example	10
Description	How often (in seconds) to perform the liveness probe
Key	pxc.livenessProbes.successThreshold
Value	int
Example	1
Description	Minimum consecutive successes for the liveness probe to be considered successful after having failed
Key	pxc.livenessProbes.failureThreshold
Value	int
Example	3
Description	When the liveness probe fails, Kubernetes will try this number of times before restarting the container
Key	pxc.envVarsSecret
Value	string
Example	my-env-var-secrets
Description	A secret with environment variables, see Define environment variables for details
Key	pxc.resources.requests.memory
Value	string
Example	1G
Description	The Kubernetes memory requests for a Percona XtraDB Cluster container
Key	pxc.resources.requests.cpu
Value	string
Example	600m
Description	Kubernetes CPU requests for a Percona XtraDB Cluster container
Key	pxc.resources.requests.ephemeral-storage
Value	string

Example	16
Description	Kubernetes Ephemeral Storage requests for a Percona XtraDB Cluster container
Key	pxc.resources.limits.memory
Value	string
Example	16
Description	Kubernetes memory limits for a Percona XtraDB Cluster container
Key	pxc.resources.limits.cpu
Value	string
Example	1
Description	Kubernetes CPU limits for a Percona XtraDB Cluster container
Key	pxc.resources.limits.ephemeral-storage
Value	string
Example	16
Description	Kubernetes Ephemeral Storage limits for a Percona XtraDB Cluster container
Key	pxc.nodeSelector
Value	label
Example	disktype: <code>ssd</code>
Description	Kubernetes nodeSelector
Key	pxc.affinity.topologyKey
Value	string
Example	<code>kubernetes.io/hostname</code>
Description	The Operator topology key node anti-affinity constraint
Key	pxc.affinity.advanced
Value	subdoc
Example	
Description	In cases where the Pods require complex tuning the advanced option turns off the <code>topologyKey</code> effect. This setting allows the standard Kubernetes affinity constraints of any complexity to be used
Key	pxc.tolerations
Value	subdoc
Example	<code>node.alpha.kubernetes.io/unreachable</code>
Description	Kubernetes Pod tolerations
Key	pxc.podDisruptionBudget.maxUnavailable
Example	1

Description	The Kubernetes podDisruptionBudget specifies the number of Pods from the set unavailable after the eviction
Key	<code>pxc.podDisruptionBudget.minAvailable</code>
Value	int
Example	0
Description	The Kubernetes podDisruptionBudget Pods that must be available after an eviction
Key	<code>pxc.volumeSpec.emptyDir</code>
Value	string
Example	{}
Description	The Kubernetes emptyDir volume The directory created on a node and accessible to the Percona XtraDB Cluster Pod containers
Key	<code>pxc.volumeSpec.hostPath.path</code>
Value	string
Example	/data
Description	Kubernetes hostPath The volume that mounts a directory from the host node's filesystem into your Pod. The path property is required
Key	<code>pxc.volumeSpec.hostPath.type</code>
Value	string
Example	Directory
Description	The Kubernetes hostPath . An optional property for the hostPath
Key	<code>pxc.volumeSpec.persistentVolumeClaim.storageClassName</code>
Value	string
Example	standard
Description	Set the Kubernetes storage class to use with the Percona XtraDB Cluster PersistentVolumeClaim
Key	<code>pxc.volumeSpec.persistentVolumeClaim.accessModes</code>
Value	array
Example	[ReadWriteOnce]
Description	The Kubernetes PersistentVolumeClaim access modes for the Percona XtraDB cluster
Key	<code>pxc.volumeSpec.resources.requests.storage</code>
Value	string
Example	6Gi
Description	The Kubernetes PersistentVolumeClaim size for the Percona XtraDB cluster
Key	<code>pxc.gracePeriod</code>
Value	int

Example	600
Description	The Kubernetes grace period when terminating a Pod
Key	pxc.containerSecurityContext
Value	subdoc
Example	<code>privileged: true</code>
Description	A custom Kubernetes Security Context for a Container to be used instead of the default one
Key	pxc.podSecurityContext
Value	subdoc
Example	<code>fsGroup: 1001 supplementalGroups: [1001, 1002, 1003]</code>
Description	A custom Kubernetes Security Context for a Pod to be used instead of the default one
Key	pxc.serviceAccountName
Value	string
Example	<code>percona-xtradb-cluster-operator-workload</code>
Description	The Kubernetes Service Account for Percona XtraDB Cluster Pods
Key	pxc.imagePullPolicy
Value	string
Example	<code>Always</code>
Description	The policy used to update images
Key	pxc.runtimeClassName
Value	string
Example	<code>image-rc</code>
Description	Name of the Kubernetes Runtime Class for Percona XtraDB Cluster Pods
Key	pxc.sidecars.image
Value	string
Example	<code>busybox</code>
Description	Image for the custom sidecar container for Percona XtraDB Cluster Pods
Key	pxc.sidecars.command
Value	array
Example	<code>["/bin/sh"]</code>
Description	Command for the custom sidecar container for Percona XtraDB Cluster Pods
Key	pxc.sidecars.args
Value	array
Example	<code>["-c", "while true; do trap 'exit 0' SIGINT SIGTERM SIGQUIT SIGKILL; done;"]</code>

Description	Command arguments for the custom sidecar container for Percona XtraDB Cluster Pods
Key	pxc.sidecars.name
Value	string
Example	my-sidecar-1
Description	Name of the custom sidecar container for Percona XtraDB Cluster Pods
Key	pxc.sidecars.resources.requests.memory
Value	string
Example	1G
Description	The Kubernetes memory requests for a Percona XtraDB Cluster sidecar container
Key	pxc.sidecars.resources.requests.cpu
Value	string
Example	500m
Description	Kubernetes CPU requests for a Percona XtraDB Cluster sidecar container
Key	pxc.sidecars.resources.limits.memory
Value	string
Example	2G
Description	Kubernetes memory limits for a Percona XtraDB Cluster sidecar container
Key	pxc.sidecars.resources.limits.cpu
Value	string
Example	600m
Description	Kubernetes CPU limits for a Percona XtraDB Cluster sidecar container

15.1.4 HAProxy section

The `haproxy` section in the [deploy/cr.yaml](#) file contains configuration options for the HAProxy service.

Key	haproxy.enabled
Value	boolean
Example	<code>true</code>
Description	Enables or disables load balancing with HAProxy Services
Key	haproxy.size
Value	int
Example	<code>2</code>
Description	The number of the HAProxy Pods to provide load balancing . It should be 2 or more unless the <code>spec.allowUnsafeConfigurations</code> key is set to true
Key	haproxy.image
Value	string
Example	<code>percona/percona-xtradb-cluster-operator:1.12.0-haproxy</code>
Description	HAProxy Docker image to use
Key	haproxy.imagePullPolicy
Value	string
Example	<code>Always</code>
Description	The policy used to update images
Key	haproxy.imagePullSecrets.name
Value	string
Example	<code>private-registry-credentials</code>
Description	The Kubernetes imagePullSecrets for the HAProxy image
Key	haproxy.readinessDelaySec
Value	int
Example	<code>15</code>
Description	Adds a delay before a run check to verify the application is ready to process traffic
Key	haproxy.livenessDelaySec
Value	int
Example	<code>300</code>
Description	Adds a delay before the run check ensures the application is healthy and capable of processing requests
Key	haproxy.configuration
Value	string
Example	
Description	The custom HAProxy configuration file contents

Key	haproxy.annotations
Value	label
Example	<code>iam.amazonaws.com/role: role-arn</code>
Description	The Kubernetes annotations metadata
Key	haproxy.labels
Value	label
Example	<code>rack: rack-22</code>
Description	Labels are key-value pairs attached to objects
Key	haproxy.readinessProbes.initialDelaySeconds
Value	int
Example	<code>15</code>
Description	Number of seconds to wait before performing the first readiness probe
Key	haproxy.readinessProbes.timeoutSeconds
Value	int
Example	<code>1</code>
Description	Number of seconds after the container has started before readiness probes are initiated
Key	haproxy.readinessProbes.periodSeconds
Value	int
Example	<code>5</code>
Description	How often (in seconds) to perform the readiness probe
Key	haproxy.readinessProbes.successThreshold
Value	int
Example	<code>1</code>
Description	Minimum consecutive successes for the readiness probe to be considered successful after having failed
Key	haproxy.readinessProbes.failureThreshold
Value	int
Example	<code>3</code>
Description	When the readiness probe fails, Kubernetes will try this number of times before marking the Pod Unready
Key	haproxy.serviceType
Value	string
Example	<code>ClusterIP</code>
Description	Specifies the type of Kubernetes Service to be used for HAProxy
Key	haproxy.externalTrafficPolicy

Value	string
Example	Cluster
Description	Specifies whether Service for HAProxy should route external traffic to cluster-wide or to node-local endpoints (it can influence the load balancing effectiveness)
Key	haproxy.resources.requests.memory
Value	string
Example	1G
Description	The Kubernetes memory requests for the main HAProxy container
Key	haproxy.resources.requests.cpu
Value	string
Example	600m
Description	Kubernetes CPU requests for the main HAProxy container
Key	haproxy.resources.limits.memory
Value	string
Example	1G
Description	Kubernetes memory limits for the main HAProxy container
Key	haproxy.resources.limits.cpu
Value	string
Example	700m
Description	Kubernetes CPU limits for the main HAProxy container
Key	haproxy.envVarsSecret
Value	string
Example	my-env-var-secrets
Description	A secret with environment variables, see Define environment variables for details
Key	haproxy.priorityClassName
Value	string
Example	high-priority
Description	The Kubernetes Pod Priority class for HAProxy
Key	haproxy.schedulerName
Value	string
Example	mycustom-scheduler
Description	The Kubernetes Scheduler
Key	haproxy.nodeSelector

Value	label
Example	disktype: ssd
Description	Kubernetes nodeSelector
Key	haproxy.affinity.topologyKey
Value	string
Example	kubernetes.io/hostname
Description	The Operator topology key node anti-affinity constraint
Key	haproxy.affinity.advanced
Value	subdoc
Example	
Description	If available it makes a topologyKey node affinity constraint to be ignored
Key	haproxy.tolerations
Value	subdoc
Example	node.alpha.kubernetes.io/unreachable
Description	Kubernetes Pod tolerations
Key	haproxy.podDisruptionBudget.maxUnavailable
Value	int
Example	1
Description	The Kubernetes podDisruptionBudget specifies the number of Pods from the set unavailable after the eviction
Key	haproxy.podDisruptionBudget.minAvailable
Value	int
Example	0
Description	The Kubernetes podDisruptionBudget Pods that must be available after an eviction
Key	haproxy.gracePeriod
Value	int
Example	30
Description	The Kubernetes grace period when terminating a Pod
Key	haproxy.loadBalancerSourceRanges
Value	string
Example	10.0.0.0/8
Description	The range of client IP addresses from which the load balancer should be reachable (if not set, there is no limitations)
Key	haproxy.loadBalancerIP

Value	string
Example	127.0.0.1
Description	The static IP-address for the load balancer
Key	haproxy.serviceLabels
Value	label
Example	rack: rack-23
Description	The Kubernetes labels for the load balancer Service
Key	haproxy.serviceAnnotations
Value	string
Example	service.beta.kubernetes.io/aws-load-balancer-backend-protocol: http
Description	The Kubernetes annotations metadata for the load balancer Service
Key	haproxy.replicasServiceEnabled
Value	boolean
Example	true
Description	Enables or disables <code>haproxy-replicas</code> Service. This Service (on by default) forwards requests to all Percona XtraDB Cluster instances, and it should not be used for write requests!
Key	haproxy.replicasLoadBalancerSourceRanges
Value	string
Example	10.0.0.0/8
Description	The range of client IP addresses from which the load balancer should be reachable (if not set, there is no limitations)
Key	haproxy.replicasLoadBalancerIP
Value	string
Example	127.0.0.1
Description	The static IP-address for the replicas load balancer
Key	haproxy.replicasServiceType
Value	string
Example	ClusterIP
Description	Specifies the type of Kubernetes Service to be used for HAProxy replicas
Key	haproxy.replicasExternalTrafficPolicy
Value	string
Example	Cluster
Description	Specifies whether Service for HAProxy replicas should route external traffic to cluster-wide or to node-local endpoints (it can influence the load balancing effectiveness)

Key	haproxy.replicasServiceLabels
Value	label
Example	<code>rack: rack-23</code>
Description	The Kubernetes labels for the <code>haproxy-replicas</code> Service
Key	haproxy.replicasServiceAnnotations
Value	string
Example	<code>service.beta.kubernetes.io/aws-load-balancer-backend-protocol: http</code>
Description	The Kubernetes annotations metadata for the <code>haproxy-replicas</code> Service
Key	haproxy.containerSecurityContext
Value	subdoc
Example	<code>privileged: true</code>
Description	A custom Kubernetes Security Context for a Container to be used instead of the default one
Key	haproxy.podSecurityContext
Value	subdoc
Example	<code>fsGroup: 1001</code> <code>supplementalGroups: [1001, 1002, 1003]</code>
Description	A custom Kubernetes Security Context for a Pod to be used instead of the default one
Key	haproxy.serviceAccountName
Value	string
Example	<code>percona-xtradb-cluster-operator-workload</code>
Description	The Kubernetes Service Account for the HAProxy Pod
Key	haproxy.runtimeClassName
Value	string
Example	<code>image-rc</code>
Description	Name of the Kubernetes Runtime Class for the HAProxy Pod
Key	haproxy.sidecars.image
Value	string
Example	<code>busybox</code>
Description	Image for the custom sidecar container for the HAProxy Pod
Key	haproxy.sidecars.command
Value	array
Example	<code>["/bin/sh"]</code>
Description	Command for the custom sidecar container for the HAProxy Pod
Key	haproxy.sidecars.args

Value	array
Example	<code>["-c", "while true; do trap 'exit 0' SIGINT SIGTERM SIGQUIT SIGKILL; done;"]</code>
Description	Command arguments for the custom sidecar container for the HAProxy Pod
Key	haproxy.sidecars.name
Value	string
Example	<code>my-sidecar-1</code>
Description	Name of the custom sidecar container for the HAProxy Pod
Key	haproxy.sidecars.resources.requests.memory
Value	string
Example	<code>16</code>
Description	The Kubernetes memory requests for the sidecar HAProxy containers
Key	haproxy.sidecars.resources.requests.cpu
Value	string
Example	<code>500m</code>
Description	Kubernetes CPU requests for the sidecar HAProxy containers
Key	haproxy.sidecars.resources.limits.memory
Value	string
Example	<code>2G</code>
Description	Kubernetes memory limits for the sidecar HAProxy containers
Key	haproxy.sidecars.resources.limits.cpu
Value	string
Example	<code>600m</code>
Description	Kubernetes CPU limits for the sidecar HAProxy containers

15.1.5 ProxySQL section

The `proxysql` section in the [deploy/cr.yaml](#) file contains configuration options for the ProxySQL daemon.

Key	<code>proxysql.enabled</code>
Value	boolean
Example	<code>false</code>
Description	Enables or disables load balancing with ProxySQL Services
Key	<code>proxysql.size</code>
Value	int
Example	<code>2</code>
Description	The number of the ProxySQL daemons to provide load balancing . It should be 2 or more unless the <code>spec.allowUnsafeConfigurations</code> key is set to true
Key	<code>proxysql.image</code>
Value	string
Example	<code>percona/percona-xtradb-cluster-operator:1.12.0-proxysql</code>
Description	ProxySQL Docker image to use
Key	<code>proxysql.imagePullPolicy</code>
Value	string
Example	<code>Always</code>
Description	The policy used to update images
Key	<code>proxysql.imagePullSecrets.name</code>
Value	string
Example	<code>private-registry-credentials</code>
Description	The Kubernetes imagePullSecrets for the ProxySQL image
Key	<code>proxysql.readinessDelaySec</code>
Value	int
Example	<code>15</code>
Description	Adds a delay before a run check to verify the application is ready to process traffic
Key	<code>proxysql.livenessDelaySec</code>
Value	int
Example	<code>300</code>
Description	Adds a delay before the run check ensures the application is healthy and capable of processing requests
Key	<code>proxysql.configuration</code>
Value	string
Example	
Description	The custom ProxySQL configuration file contents

Key	proxysql.annotations
Value	label
Example	<code>iam.amazonaws.com/role: role-arn</code>
Description	The Kubernetes annotations metadata
Key	proxysql.labels
Value	label
Example	<code>rack: rack-22</code>
Description	Labels are key-value pairs attached to objects
Key	proxysql.serviceType
Value	string
Example	<code>ClusterIP</code>
Description	Specifies the type of Kubernetes Service to be used
Key	proxysql.externalTrafficPolicy
Value	string
Example	<code>Cluster</code>
Description	Specifies whether Service should route external traffic to cluster-wide or node-local endpoints (it can influence the load balancing effectiveness)
Key	proxysql.resources.requests.memory
Value	string
Example	<code>1G</code>
Description	The Kubernetes memory requests for the main ProxySQL container
Key	proxysql.resources.requests.cpu
Value	string
Example	<code>600m</code>
Description	Kubernetes CPU requests for the main ProxySQL container
Key	proxysql.resources.limits.memory
Value	string
Example	<code>1G</code>
Description	Kubernetes memory limits for the main ProxySQL container
Key	proxysql.resources.limits.cpu
Value	string
Example	<code>700m</code>
Description	Kubernetes CPU limits for the main ProxySQL container

Key	proxysql.envVarsSecret
Value	string
Example	<code>my-env-var-secrets</code>
Description	A secret with environment variables, see Define environment variables for details
Key	proxysql.priorityClassName
Value	string
Example	<code>high-priority</code>
Description	The Kubernetes Pod Priority class for ProxySQL
Key	proxysql.schedulerName
Value	string
Example	<code>mycustom-scheduler</code>
Description	The Kubernetes Scheduler
Key	proxysql.nodeSelector
Value	label
Example	<code>disktype: ssd</code>
Description	Kubernetes nodeSelector
Key	proxysql.affinity.topologyKey
Value	string
Example	<code>kubernetes.io/hostname</code>
Description	The Operator topology key node anti-affinity constraint
Key	proxysql.affinity.advanced
Value	subdoc
Example	
Description	If available it makes a topologyKey node affinity constraint to be ignored
Key	proxysql.tolerations
Value	subdoc
Example	<code>node.alpha.kubernetes.io/unreachable</code>
Description	Kubernetes Pod tolerations
Key	proxysql.volumeSpec.emptyDir
Value	string
Example	<code>{}</code>
Description	The Kubernetes emptyDir volume The directory created on a node and accessible to the Percona XtraDB Cluster Pod containers

Key	<code>proxysql.volumeSpec.hostPath.path</code>
Value	string
Example	<code>/data</code>
Description	Kubernetes hostPath The volume that mounts a directory from the host node's filesystem into your Pod. The path property is required
Key	<code>proxysql.volumeSpec.hostPath.type</code>
Value	string
Example	<code>Directory</code>
Description	The Kubernetes hostPath . An optional property for the hostPath
Key	<code>proxysql.volumeSpec.persistentVolumeClaim.storageClassName</code>
Value	string
Example	<code>standard</code>
Description	Set the Kubernetes storage class to use with the Percona XtraDB Cluster PersistentVolumeClaim
Key	<code>proxysql.volumeSpec.persistentVolumeClaim.accessModes</code>
Value	array
Example	<code>[ReadWriteOnce]</code>
Description	The Kubernetes PersistentVolumeClaim access modes for the Percona XtraDB cluster
Key	<code>proxysql.volumeSpec.resources.requests.storage</code>
Value	string
Example	<code>6Gi</code>
Description	The Kubernetes PersistentVolumeClaim size for the Percona XtraDB cluster
Key	<code>proxysql.podDisruptionBudget.maxUnavailable</code>
Value	int
Example	<code>1</code>
Description	The Kubernetes podDisruptionBudget specifies the number of Pods from the set unavailable after the eviction
Key	<code>proxysql.podDisruptionBudget.minAvailable</code>
Value	int
Example	<code>0</code>
Description	The Kubernetes podDisruptionBudget Pods that must be available after an eviction
Key	<code>proxysql.gracePeriod</code>
Value	int
Example	<code>30</code>
Description	The Kubernetes grace period when terminating a Pod

Key	proxysql.loadBalancerSourceRanges
Value	string
Example	<code>10.0.0.0/8</code>
Description	The range of client IP addresses from which the load balancer should be reachable (if not set, there is no limitations)
Key	proxysql.serviceLabels
Value	label
Example	<code>rack: rack-23</code>
Description	The Kubernetes labels for the load balancer Service
Key	proxysql.serviceAnnotations
Value	string
Example	<code>service.beta.kubernetes.io/aws-load-balancer-backend-protocol: http</code>
Description	The Kubernetes annotations metadata for the load balancer Service
Key	proxysql.containerSecurityContext
Value	subdoc
Example	<code>privileged: true</code>
Description	A custom Kubernetes Security Context for a Container to be used instead of the default one
Key	proxysql.podSecurityContext
Value	subdoc
Example	<code>fsGroup: 1001 supplementalGroups: [1001, 1002, 1003]</code>
Description	A custom Kubernetes Security Context for a Pod to be used instead of the default one
Key	proxysql.serviceAccountName
Value	string
Example	<code>percona-xtradb-cluster-operator-workload</code>
Description	The Kubernetes Service Account for the ProxySQL Pod
Key	proxysql.runtimeClassName
Value	string
Example	<code>image-rc</code>
Description	Name of the Kubernetes Runtime Class for the ProxySQL Pod
Key	proxysql.sidecars.image
Value	string
Example	<code>busybox</code>
Description	Image for the custom sidecar container for the ProxySQL Pod

Key	<code>proxysql.sidecars.command</code>
Value	array
Example	<code>["/bin/sh"]</code>
Description	Command for the custom sidecar container for the ProxySQL Pod
Key	<code>proxysql.sidecars.args</code>
Value	array
Example	<code>["-c", "while true; do trap 'exit 0' SIGINT SIGTERM SIGQUIT SIGKILL; done;"]</code>
Description	Command arguments for the custom sidecar container for the ProxySQL Pod
Key	<code>proxysql.sidecars.name</code>
Value	string
Example	<code>my-sidecar-1</code>
Description	Name of the custom sidecar container for the ProxySQL Pod
Key	<code>proxysql.sidecars.resources.requests.memory</code>
Value	string
Example	<code>16</code>
Description	The Kubernetes memory requests for the sidecar ProxySQL containers
Key	<code>proxysql.sidecars.resources.requests.cpu</code>
Value	string
Example	<code>500m</code>
Description	Kubernetes CPU requests for the sidecar ProxySQL containers
Key	<code>proxysql.sidecars.resources.limits.memory</code>
Value	string
Example	<code>26</code>
Description	Kubernetes memory limits for the sidecar ProxySQL containers
Key	<code>proxysql.sidecars.resources.limits.cpu</code>
Value	string
Example	<code>600m</code>
Description	Kubernetes CPU limits for the sidecar ProxySQL containers

15.1.6 Log Collector section

The `logcollector` section in the `deploy/cr.yaml` file contains configuration options for [Fluent Bit Log Collector](#).

Key	<code>logcollector.enabled</code>
Value	boolean
Example	<code>true</code>
Description	Enables or disables cluster-level logging with Fluent Bit
Key	<code>logcollector.image</code>
Value	string
Example	<code>percona/percona-xtradb-cluster-operator:1.6.0-logcollector</code>
Description	Log Collector Docker image to use
Key	<code>logcollector.configuration</code>
Value	subdoc
Example	
Description	Additional configuration options (see Fluent Bit official documentation for details)
Key	<code>logcollector.resources.requests.memory</code>
Value	string
Example	<code>100M</code>
Description	The Kubernetes memory requests for a Log Collector container
Key	<code>logcollector.resources.requests.cpu</code>
Value	string
Example	<code>200m</code>
Description	Kubernetes CPU requests for a Log collector container

15.1.7 PMM section

The `pmm` section in the [deploy/cr.yaml](#) file contains configuration options for Percona Monitoring and Management.

Key	pmm.enabled
Value	boolean
Example	false
Description	Enables or disables monitoring Percona XtraDB cluster with PMM
Key	pmm.image
Value	string
Example	percona/pmm-client:2.32.0
Description	PMM client Docker image to use
Key	pmm.serverHost
Value	string
Example	monitoring-service
Description	Address of the PMM Server to collect data from the cluster
Key	pmm.serverUser
Value	string
Example	admin
Description	The PMM Serve_User . The PMM Server password should be configured using Secrets
Key	pmm.resources.requests.memory
Value	string
Example	150M
Description	The Kubernetes memory requests for a PMM container
Key	pmm.resources.requests.cpu
Value	string
Example	300m
Description	Kubernetes CPU requests for a PMM container
Key	pmm.pxcParams
Value	string
Example	--disable-tablestats-limit=2000
Description	Additional parameters which will be passed to the pmm-admin add mysql command for <code>pxc</code> Pods
Key	pmm.proxysqlParams
Value	string
Example	--custom-labels=CUSTOM-LABELS
Description	Additional parameters which will be passed to the pmm-admin add mysql command for <code>proxysql</code> Pods

15.1.8 Backup section

The `backup` section in the [deploy/cr.yaml](#) file contains the following configuration options for the regular Percona XtraDB Cluster backups.

Key	<code>backup.image</code>
Value	string
Example	<code>percona/percona-xtradb-cluster-operator:1.12.0-backup</code>
Description	The Percona XtraDB cluster Docker image to use for the backup
Key	<code>backup.backoffLimit</code>
Value	int
Example	6
Description	The number of retries to make a backup
Key	<code>backup.imagePullSecrets.name</code>
Value	string
Example	<code>private-registry-credentials</code>
Description	The Kubernetes imagePullSecrets for the specified image
Key	<code>backup.storages.<storage-name>.type</code>
Value	string
Example	<code>s3</code>
Description	The cloud storage type used for backups. Only <code>s3</code> and <code>filesystem</code> types are supported
Key	<code>backup.storages.<storage-name>.verifyTLS</code>
Value	boolean
Example	<code>true</code>
Description	Enable or disable verification of the storage server TLS certificate. Disabling it may be useful e.g. to skip TLS verification for private S3-compatible storage with a self-issued certificate
Key	<code>backup.storages.<storage-name>.s3.credentialsSecret</code>
Value	string
Example	<code>my-cluster-name-backup-s3</code>
Description	The Kubernetes secret for backups. It should contain <code>AWS_ACCESS_KEY_ID</code> and <code>AWS_SECRET_ACCESS_KEY</code> keys
Key	<code>backup.storages.<storage-name>.s3.bucket</code>
Value	string
Example	
Description	The Amazon S3 bucket name for backups
Key	<code>backup.storages.s3.<storage-name>.region</code>
Value	string
Example	<code>us-east-1</code>
Description	The AWS region to use. Please note this option is mandatory for Amazon and all S3-compatible storages

Key	<code>backup.storages.s3.<storage-name>.endpointUrl</code>
Value	string
Example	
Description	The endpoint URL of the S3-compatible storage to be used (not needed for the original Amazon S3 cloud)
Key	<code>backup.storages.<storage-name>.persistentVolumeClaim.type</code>
Value	string
Example	<code>filesystem</code>
Description	The persistent volume claim storage type
Key	<code>backup.storages.<storage-name>.persistentVolumeClaim.storageClassName</code>
Value	string
Example	<code>standard</code>
Description	Set the Kubernetes Storage Class to use with the Percona XtraDB Cluster backups PersistentVolumeClaims for the <code>filesystem</code> storage type
Key	<code>backup.storages.<storage-name>.volume.persistentVolumeClaim.accessModes</code>
Value	array
Example	<code>[ReadWriteOnce]</code>
Description	The Kubernetes PersistentVolume access modes
Key	<code>backup.storages.<storage-name>.volume.persistentVolumeClaim.resources.requests.storage</code>
Value	string
Example	<code>6Gi</code>
Description	Storage size for the PersistentVolume
Key	<code>backup.storages.<storage-name>.annotations</code>
Value	label
Example	<code>iam.amazonaws.com/role: role-arn</code>
Description	The Kubernetes annotations
Key	<code>backup.storages.<storage-name>.labels</code>
Value	label
Example	<code>rack: rack-22</code>
Description	Labels are key-value pairs attached to objects
Key	<code>backup.storages.<storage-name>.resources.requests.memory</code>
Value	string
Example	<code>1G</code>
Description	The Kubernetes memory requests for a Percona XtraDB Cluster container

Key	<code>backup.storages.<storage-name>.resources.requests.cpu</code>
Value	string
Example	600m
Description	Kubernetes CPU requests for a Percona XtraDB Cluster container
Key	<code>backup.storages.<storage-name>.resources.limits.memory</code>
Value	string
Example	1G
Description	Kubernetes memory limits for a Percona XtraDB Cluster container
Key	<code>backup.storages.<storage-name>.nodeSelector</code>
Value	label
Example	disktype: ssd
Description	Kubernetes nodeSelector
Key	<code>backup.storages.<storage-name>.affinity.nodeAffinity</code>
Value	subdoc
Example	
Description	The Operator node affinity constraint
Key	<code>backup.storages.<storage-name>.tolerations</code>
Value	subdoc
Example	<code>backupWorker</code>
Description	Kubernetes Pod tolerations
Key	<code>backup.storages.<storage-name>.priorityClassName</code>
Value	string
Example	<code>high-priority</code>
Description	The Kubernetes Pod priority class
Key	<code>backup.storages.<storage-name>.schedulerName</code>
Value	string
Example	<code>mycustom-scheduler</code>
Description	The Kubernetes Scheduler
Key	<code>backup.storages.<storage-name>.containerSecurityContext</code>
Value	subdoc
Example	<code>privileged: true</code>
Description	A custom Kubernetes Security Context for a Container to be used instead of the default one
Key	<code>backup.storages.<storage-name>.podSecurityContext</code>

Value	subdoc
Example	fsGroup: 1001 supplementalGroups: [1001, 1002, 1003]
Description	A custom Kubernetes Security Context for a Pod to be used instead of the default one
Key	backup.schedule.name
Value	string
Example	sat-night-backup
Description	The backup name
Key	backup.schedule.schedule
Value	string
Example	0 0 * * 6
Description	Scheduled time to make a backup specified in the crontab format
Key	backup.schedule.keep
Value	int
Example	3
Description	The amount of most recent backups to store. Older backups are automatically deleted. Set <code>keep</code> to zero or completely remove it to disable automatic deletion of backups
Key	backup.schedule.storageName
Value	string
Example	s3-us-west
Description	The name of the storage for the backups configured in the <code>storages</code> or <code>fs-pvc</code> subsection
Key	backup.pitr.enabled
Value	boolean
Example	false
Description	Enables or disables point-in-time-recovery functionality
Key	backup.pitr.storageName
Value	string
Example	s3-us-west
Description	The name of the storage for the backups configured in the <code>storages</code> subsection, which will be reused to store binlog for point-in-time-recovery
Key	backup.pitr.timeBetweenUploads
Value	int
Example	60
Description	Seconds between running the binlog uploader

Last update: 2022-12-07

15.2 Percona certified images

Following table presents Percona's certified docker images to be used with the Percona Operator for MySQL based on Percona XtraDB Cluster:

Image	Digest
percona/percona-xtradb-cluster-operator: 1.12.0	fbbf093bf7681e0784dd817ac311dea815140ce53aa9cdc79ada0bba067b06a8
percona/percona-xtradb-cluster-operator: 1.12.0-haproxy	d900211bf5684839cfbaab3ec939ef7bae770638c7d819904820a3882a2aea32
percona/percona-xtradb-cluster-operator: 1.12.0-proxysql	a0352f7491ba32f4cb1fb8edb1ad9248235f4ae1186932965383f47e3239e0f7
percona/percona-xtradb-cluster-operator: 1.12.0-pxc8.0-backup	d5a45266857dcccec973103cd5ad35306bdd83d540fa1f31be673aa12814d521
percona/percona-xtradb-cluster-operator: 1.12.0-pxc5.7-backup	5ec440d514f57a5a49e62849674bd4df33ff3e7697037ef4b7ead9dd346af5f9
percona/percona-xtradb-cluster-operator: 1.12.0-logcollector	6d4a4537fa766743c5aa6da8b683bf2656980d6979c31eab50a4be69c80f51ce
percona/pmm-client:2.32.0	ee2f3db541857e0a71633270596933441c4be579ce8e33c22cf150ead4f3622f
percona/percona-xtradb-cluster:8.0.29-21.1	96c6bb8189280aeb773e74ed46aa41c01781b62947ed70c89efeb9f41c367ee7
percona/percona-xtradb-cluster:8.0.27-18.1	a0fced75ecd2cd164dd9937917440911aed972476d48a2b8a84fe832bc67e43a
percona/percona-xtradb-cluster:8.0.25-15.1	529e979c86442429e6feabef9a2d9fc362f4626146f208fbfac704e145a492dd
percona/percona-xtradb-cluster:8.0.23-14.1	8109f7ca4fc465ba862c08021df12e77b65d384395078e31e270d14b77810d79
percona/percona-xtradb-cluster:8.0.22-13.1	1295af1153c1d02e9d40131eb0945b53f7f371796913e64116bf2caa77dc186d
percona/percona-xtradb-cluster:8.0.21-12.1	d95cf39a58f09759408a00b519fe0d0b19c1b28332ece94349dd5e9cddbda017e
percona/percona-xtradb-cluster:8.0.20-11.2	feda5612db18da824e971891d6084465aa9cdc9918c18001cd95ba30916da78b
percona/percona-xtradb-cluster:8.0.20-11.1	54b1b2f5153b78b05d651034d4603a13e685cbb9b45bfa09a39864fa3f169349
percona/percona-xtradb-cluster:8.0.19-10.1	1058ae8eded735ebdf664807aad7187942fc9a1170b3fd0369574cb61206b63a
percona/percona-xtradb-cluster: 5.7.39-31.61	9013170a71559bbac92ba9c2e986db9bda3a8a9e39ee1ee350e0ee94488bb6d7
percona/percona-xtradb-cluster: 5.7.36-31.55	c7bad990fc7ca0fde89240e921052f49da08b67c7c6dc54239593d61710be504
percona/percona-xtradb-cluster: 5.7.35-31.53	4cc61888821d5b96942ebd1637e5622a1bcc607f6b7e0043dee3d430460ea75
percona/percona-xtradb-cluster: 5.7.34-31.51	f8d51d7932b9bb1a5a896c7ae440256230eb69b55798ff37397aabfd58b80ccb
percona/percona-xtradb-cluster: 5.7.33-31.49	f0a4bbb0ec5adff2a2d3e88194b3dcac479266ca29da028f0dfb22f55449ac17
percona/percona-xtradb-cluster: 5.7.32-31.47	7b095019ad354c336494248d6080685022e2ed46e3b53fc103b25cd12c95952b
percona/percona-xtradb-cluster: 5.7.31-31.45.2	0defc85c7c7afacc438f5fe355dc8320ea7ffc7018ca2cb6bda3ac0c526ae172
percona/percona-xtradb-cluster: 5.7.31-31.45	3852cef43c0c6aa791463ba6279e59dcdac3a4fb1a5616c745c1b3c68041dc2
percona/percona-xtradb-cluster: 5.7.30-31.43	b03a060e9261b37288a2153c78f86dcfc53367c36e1bcdcae046dd2d0b0721af
percona/percona-xtradb-cluster: 5.7.29-31.43	85fb479de073770280ae601cf3ec22dc5c8cca4c8b0dc893b09503767338e6f9

Image	Digest
percona/percona-xtradb-cluster: 5.7.28-31.41.2	fcd6525aaeedb5e436e9534e2a63aebcf743c043526dd05dba8519ebddc8b30
percona/percona-xtradb-cluster: 5.7.27-31.39	7d8eb4d2031c32c6e96451655f359d8e5e8e047dc95bada9a28c41c158876c26
percona/percona-xtradb-cluster: 5.7.26-31.37	9d43d8e435e4aca5c694f726cc736667cb938158635c5f01a0e9412905f1327f

Last update: 2022-12-07

15.3 Percona Operator for MySQL API Documentation

Percona Operator for MySQL based on Percona XtraDB Cluster provides an [aggregation-layer extension for the Kubernetes API](#). Please refer to the [official Kubernetes API documentation](#) on the API access and usage details. The following subsections describe the Percona XtraDB Cluster API provided by the Operator.

15.3.1 Prerequisites

1. Create the namespace name you will use, if not exist:

```
$ kubectl create namespace my-namespace-name
```

Trying to create an already-existing namespace will show you a self-explanatory error message. Also, you can use the `default` namespace.

Note

In this document `default` namespace is used in all examples. Substitute `default` with your namespace name if you use a different one.

2. Prepare

```
# set correct API address
KUBE_CLUSTER=$(kubectl config view --minify -o jsonpath='{.clusters[0].name}')
API_SERVER=$(kubectl config view -o jsonpath="{.clusters[?(@.name=\"$KUBE_CLUSTER\")].cluster.server}" | sed -e 's#https://##')

# create service account and get token
kubectl apply -f deploy/crd.yaml -f deploy/rbac.yaml -n default
KUBE_TOKEN=$(kubectl get secret $(kubectl get serviceaccount percona-xtradb-cluster-operator -o jsonpath='{.secrets[0].name}' -n default) -o jsonpath='{.data.token}' -n default | base64 --decode)
```

15.3.2 Create new Percona XtraDB Cluster

Description:

The `command` to create a new Percona XtraDB Cluster with all its resources

Kubectl Command:

```
$ kubectl apply -f percona-xtradb-cluster-operator/deploy/cr.yaml
```

URL:

```
https://$API_SERVER/apis/pxc.percona.com/v1-12-0/namespaces/default/perconaxtradbclusters
```

Authentication:

```
Authorization: Bearer $KUBE_TOKEN
```

cURL Request:

```
$ curl -k -v -XPOST "https://$API_SERVER/apis/pxc.percona.com/v1-12-0/namespaces/default/perconaxtradbclusters" \
-H "Content-Type: application/json" \
-H "Accept: application/json" \
-H "Authorization: Bearer $KUBE_TOKEN" \
-d "@cluster.json"
```

Request Body (cluster.json):

Example

```
{
  "apiVersion": "pxc.percona.com/v1-5-0",
  "kind": "PerconaXtraDBCluster",
  "metadata": {
    "name": "cluster1",
    "finalizers": [
      "delete-pxc-pods-in-order"
    ]
  },
  "spec": {
    "secretsName": "my-cluster-secrets",
    "vaultSecretName": "keyring-secret-vault",
    "sslSecretName": "my-cluster-ssl",
    "sslInternalSecretName": "my-cluster-ssl-internal",
    "allowUnsafeConfigurations": true,
    "pxc": {
      "size": 3,
      "image": "percona/percona-xtradb-cluster:8.0.19-10.1",
      "resources": {
        "requests": null
      },
      "affinity": {
        "antiAffinityTopologyKey": "none"
      },
      "podDisruptionBudget": {
        "maxUnavailable": 1
      },
      "volumeSpec": {
        "persistentVolumeClaim": {
          "resources": {
            "requests": {
              "storage": "6Gi"
            }
          }
        }
      },
      "gracePeriod": 600
    },
    "proxysql": {
      "enabled": true,
      "size": 3,
      "image": "percona/percona-xtradb-cluster-operator:1.5.0-proxysql",
      "resources": {
        "requests": null
      },
      "affinity": {
        "antiAffinityTopologyKey": "none"
      },
      "volumeSpec": {
        "persistentVolumeClaim": {
          "resources": {
            "requests": {
              "storage": "2Gi"
            }
          }
        }
      },
      "podDisruptionBudget": {
        "maxUnavailable": 1
      },
      "gracePeriod": 30
    },
    "pmm": {
      "enabled": false,
      "image": "percona/percona-xtradb-cluster-operator:1.5.0-pmm",
      "serverHost": "monitoring-service",
      "serverUser": "pmm"
    },
    "backup": {
      "image": "percona/percona-xtradb-cluster-operator:1.5.0-pxc8.0-backup",
      "serviceAccountName": "percona-xtradb-cluster-operator",
      "storages": {
        "s3-us-west": {
          "type": "s3",
          "s3": {
            "bucket": "S3-BACKUP-BUCKET-NAME-HERE",
            "credentialsSecret": "my-cluster-name-backup-s3",
            "region": "us-west-2"
          }
        },
        "fs-pvc": {
          "type": "filesystem",
          "volume": {
            "persistentVolumeClaim": {
              "accessModes": [
                "ReadWriteOnce"
              ],
              "resources": {
                "requests": {
                  "storage": "6Gi"
                }
              }
            }
          }
        }
      }
    }
  }
}
```

```

    "schedule": [
      {
        "name": "sat-night-backup",
        "schedule": "0 0 * * 6",
        "keep": 3,
        "storageName": "s3-us-west"
      },
      {
        "name": "daily-backup",
        "schedule": "0 0 * * *",
        "keep": 5,
        "storageName": "fs-pvc"
      }
    ]
  }
}

```

Inputs:**Metadata:**

1. Name (String, min-length: 1) : contains name of cluster
2. Finalizers (list of string, Default: ["delete-pxc-pods-in-order"]) contains steps to do when deleting the cluster

Spec:

1. secretsName (String, min-length: 1) : contains name of secret to create for the cluster
2. vaultSecretName (String, min-length: 1) : contains name of vault secret to create for the cluster
3. sslInternalSecretName (String, min-length: 1) : contains name of ssl secret to create for the cluster
4. allowUnsafeConfigurations (Boolean, Default: false) : allow unsafe configurations to run

pxc:

1. Size (Int , min-value: 1, default, 3) : number of Percona XtraDB Cluster nodes to create
2. Image (String, min-length: 1) : contains image name to use for Percona XtraDB Cluster nodes
3. volumeSpec : storage (SizeString, default: "6Gi") : contains the size for the storage volume of Percona XtraDB Cluster nodes
4. gracePeriod (Int, default: 600, min-value: 0) : contains the time to wait for Percona XtraDB Cluster node to shutdown in milliseconds

proxysql:

1. Enabled (Boolean, default: true) : enabled or disables proxysql

pmm:

1. serverHost (String, min-length: 1) : service name for monitoring
2. serverUser (String, min-length: 1) : name of pmm user
3. image (String, min-length: 1) : name of pmm image

backup:

1. Storages (Object) : contains the storage destinations to save the backups in
2. schedule:
 - a. name (String, min-length: 1) : name of backup job
 - b. schedule (String, Cron format: "* * * * *") : contains cron schedule format for when to run cron jobs
 - c. keep (Int, min-value = 1) : number of backups to keep
 - d. storageName (String, min-length: 1) : name of storage object to use

Response:

Example

```

{
  "apiVersion": "pxc.percona.com/v1-5-0",
  "kind": "PerconaXtraDBCluster",
  "metadata": {
    "creationTimestamp": "2020-05-27T22:23:58Z",
    "finalizers": [
      "delete-pxc-pods-in-order"
    ],
    "generation": 1,
    "managedFields": [
      {
        "apiVersion": "pxc.percona.com/v1-5-0",
        "fieldsType": "FieldsV1",
        "fieldsV1": {
          "f:metadata": {
            "f:finalizers": {
              }
            },
          "f:spec": {
            ".": {
              },
            "f:allowUnsafeConfigurations": {
              },
            "f:backup": {
              ".": {
                },
              "f:image": {
                },
            },
            "f:schedule": {
              },
            "f:serviceName": {
              },
            "f:storages": {
              ".": {
                },
              "f:fs-pvc": {
                ".": {
                  },
                "f:type": {
                  },
            },
            "f:volume": {
              ".": {
                },
              "f:persistentVolumeClaim": {
                ".": {
                  },
                "f:accessModes": {
                  },
            },
            "f:resources": {
              ".": {
                },
              "f:requests": {
                ".": {
                  },
                "f:storage": {
                  }
                }
              }
            },
            "f:s3-us-west": {
              ".": {
                },
              "f:s3": {
                ".": {
                  },
                "f:bucket": {
                  },
            },
            "f:credentialsSecret": {
              },
            "f:region": {
              }
            },
            "f:type": {

```

```

    }
  }
},
"f:pmi":{
  ".":{

  },
  "f:enabled":{

  },
  "f:image":{

  },
  "f:serverHost":{

  },
  "f:serverUser":{

  }
},
"f:proxysql":{
  ".":{

  },
  "f:affinity":{
    ".":{

    },
    "f:antiAffinityTopologyKey":{

    }
  },
  "f:enabled":{

  },
  "f:gracePeriod":{

  },
  "f:image":{

  },
  "f:podDisruptionBudget":{
    ".":{

    },
    "f:maxUnavailable":{

    }
  },
  "f:resources":{
    ".":{

    },
    "f:requests":{

    }
  },
  "f:size":{

  },
  "f:volumeSpec":{
    ".":{

    },
    "f:persistentVolumeClaim":{
      ".":{

      },
      "f:resources":{
        ".":{

        },
        "f:requests":{
          ".":{

          },
          "f:storage":{

          }
        }
      }
    }
  },
},
"f:pxc":{
  ".":{

  },
  "f:affinity":{
    ".":{

    },
    "f:antiAffinityTopologyKey":{

    }
  },
  "f:gracePeriod":{

  }

```

```

    },
    "f:image":{
    },
    "f:podDisruptionBudget":{
      ".":{
      },
      "f:maxUnavailable":{
      }
    },
    "f:resources":{
      ".":{
      },
      "f:requests":{
      }
    },
    "f:size":{
    },
    "f:volumeSpec":{
      ".":{
      },
      "f:persistentVolumeClaim":{
        ".":{
        },
        "f:resources":{
          ".":{
          },
          "f:requests":{
            ".":{
            },
            "f:storage":{
            }
          }
        }
      }
    },
    "f:secretsName":{
    },
    "f:sslInternalSecretName":{
    },
    "f:sslSecretName":{
    },
    "f:vaultSecretName":{
    }
  }
},
"manager":"kubectl",
"operation":"Update",
"time":"2020-05-27T22:23:58Z"
}
],
"name":"cluster1",
"namespace":"default",
"resourceVersion":"8694",
"selfLink":"/apis/pxc.percona.com/v1-5-0/namespaces/default/perconaxtradbclusters/cluster1",
"uid":"e9115e2a-49df-4ebf-9dab-fa5a550208d3"
},
"spec":{
  "allowUnsafeConfigurations":false,
  "backup":{
    "image":"percona/percona-xtradb-cluster-operator:1.5.0-pxc8.0-backup",
    "schedule":[
      {
        "keep":3,
        "name":"sat-night-backup",
        "schedule":"0 0 * * 6",
        "storageName":"s3-us-west"
      },
      {
        "keep":5,
        "name":"daily-backup",
        "schedule":"0 0 * * *",
        "storageName":"fs-pvc"
      }
    ],
    "serviceAccountName":"percona-xtradb-cluster-operator",
    "storages":{
      "fs-pvc":{
        "type":"filesystem",
        "volume":{
          "persistentVolumeClaim":{
            "accessModes":[
              "ReadWriteOnce"
            ],
            "resources":{

```

```

        "requests":{
          "storage":"6Gi"
        }
      }
    },
    "s3-us-west":{
      "s3":{
        "bucket":"S3-BACKUP-BUCKET-NAME-HERE",
        "credentialsSecret":"my-cluster-name-backup-s3",
        "region":"us-west-2"
      },
      "type":"s3"
    }
  },
  "pmm":{
    "enabled":false,
    "image":"percona/percona-xtradb-cluster-operator:1.5.0-pmm",
    "serverHost":"monitoring-service",
    "serverUser":"pmm"
  },
  "proxysql":{
    "affinity":{
      "antiAffinityTopologyKey":"none"
    },
    "enabled":true,
    "gracePeriod":30,
    "image":"percona/percona-xtradb-cluster-operator:1.5.0-proxysql",
    "podDisruptionBudget":{
      "maxUnavailable":1
    },
    "resources":{
      "requests":null
    },
    "size":3,
    "volumeSpec":{
      "persistentVolumeClaim":{
        "resources":{
          "requests":{
            "storage":"2Gi"
          }
        }
      }
    }
  },
  "pxc":{
    "affinity":{
      "antiAffinityTopologyKey":"none"
    },
    "gracePeriod":600,
    "image":"percona/percona-xtradb-cluster:8.0.19-10.1",
    "podDisruptionBudget":{
      "maxUnavailable":1
    },
    "resources":{
      "requests":null
    },
    "size":3,
    "volumeSpec":{
      "persistentVolumeClaim":{
        "resources":{
          "requests":{
            "storage":"6Gi"
          }
        }
      }
    }
  },
  "secretsName":"my-cluster-secrets",
  "sslInternalSecretName":"my-cluster-ssl-internal",
  "sslSecretName":"my-cluster-ssl",
  "vaultSecretName":"keyring-secret-vault"
}
}

```

15.3.3 List Percona XtraDB Clusters

Description:

Lists all Percona XtraDB Clusters that exist in your kubernetes cluster

Kubectl Command:

```
$ kubectl get pxc
```

URL:

[https://\\$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters?limit=500](https://$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters?limit=500)

Authentication:

```
Authorization: Bearer $KUBE_TOKEN
```

cURL Request:

```
$ curl -k -v -XGET "https://$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters?limit=500" \  
-H "Accept: application/json;as=Table;v=v1;g=meta.k8s.io,application/json;as=Table;v=v1beta1;g=meta.k8s.io,application/json" \  
-H "Authorization: Bearer $KUBE_TOKEN"
```

Request Body:

```
None
```

Response:

Example

```

{
  "kind": "Table",
  "apiVersion": "meta.k8s.io/v1",
  "metadata": {
    "selfLink": "/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters",
    "resourceVersion": "10528"
  },
  "columnDefinitions": [
    {
      "name": "Name",
      "type": "string",
      "format": "name",
      "description": "Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names",
      "priority": 0
    },
    {
      "name": "Endpoint",
      "type": "string",
      "format": "",
      "description": "Custom resource definition column (in JSONPath format): .status.host",
      "priority": 0
    },
    {
      "name": "Status",
      "type": "string",
      "format": "",
      "description": "Custom resource definition column (in JSONPath format): .status.state",
      "priority": 0
    },
    {
      "name": "PXC",
      "type": "string",
      "format": "",
      "description": "Ready pxc nodes",
      "priority": 0
    },
    {
      "name": "proxysql",
      "type": "string",
      "format": "",
      "description": "Ready pxc nodes",
      "priority": 0
    },
    {
      "name": "Age",
      "type": "date",
      "format": "",
      "description": "Custom resource definition column (in JSONPath format): .metadata.creationTimestamp",
      "priority": 0
    }
  ],
  "rows": [
    {
      "cells": [
        "cluster1",
        "cluster1-proxysql.default",
        "ready",
        "3",
        "3",
        "8m37s"
      ],
      "object": {
        "kind": "PartialObjectMetadata",
        "apiVersion": "meta.k8s.io/v1",
        "metadata": {
          "name": "cluster1",
          "namespace": "default",
          "selfLink": "/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters/cluster1",
          "uid": "e9115e2a-49df-4ebf-9dab-fa5a550208d3",
          "resourceVersion": "10517",
          "generation": 1,
          "creationTimestamp": "2020-05-27T22:23:58Z",
          "finalizers": [
            "delete-pxc-pods-in-order"
          ],
          "managedFields": [
            {
              "manager": "kubectl",
              "operation": "Update",
              "apiVersion": "pxc.percona.com/v1-5-0",
              "time": "2020-05-27T22:23:58Z",
              "fieldsType": "FieldsV1",
              "fieldsV1": {
                "f:metadata": {
                  "f:finalizers": {
                }
              }
            },
            "f:spec": {
              ".-": {
            }
          },
          "f:allowUnsafeConfigurations": {
        }
      }
    }
  ]
}

```

```

"f:backup":{
  ".":{
  },
  "f:image":{
  },
  "f:schedule":{
  },
  "f:serviceName":{
  },
  "f:storages":{
    ".":{
    },
    "f:fs-pvc":{
      ".":{
      },
      "f:type":{
      },
      "f:volume":{
        ".":{
        },
        "f:persistentVolumeClaim":{
          ".":{
          },
          "f:accessModes":{
          },
          "f:resources":{
            ".":{
            },
            "f:requests":{
              ".":{
              },
              "f:storage":{
                }
              }
            }
          }
        }
      }
    }
  },
  "f:s3-us-west":{
    ".":{
    },
    "f:s3":{
      ".":{
      },
      "f:bucket":{
      },
      "f:credentialsSecret":{
      },
      "f:region":{
      }
    }
  },
  "f:type":{
  }
}
},
"f:pmm":{
  ".":{
  },
  "f:image":{
  },
  "f:serverHost":{
  },
  "f:serverUser":{
  }
},
"f:proxysql":{
  ".":{
  },
  "f:affinity":{
    ".":{
    },
    "f:antiAffinityTopologyKey":{
    }
  }
}

```

```

    },
    "f:enabled":{
    },
    "f:gracePeriod":{
    },
    "f:image":{
    },
    "f:podDisruptionBudget":{
      ".":{
      },
      "f:maxUnavailable":{
      }
    }
  },
  "f:resources":{
  },
  "f:size":{
  },
  "f:volumeSpec":{
    ".":{
    },
    "f:persistentVolumeClaim":{
      ".":{
      },
      "f:resources":{
        ".":{
        },
        "f:requests":{
          ".":{
          },
          "f:storage":{
          }
        }
      }
    }
  }
},
"f:pxc":{
  ".":{
  },
  "f:affinity":{
    ".":{
    },
    "f:antiAffinityTopologyKey":{
    }
  },
  "f:gracePeriod":{
  },
  "f:image":{
  },
  "f:podDisruptionBudget":{
    ".":{
    },
    "f:maxUnavailable":{
    }
  },
  "f:resources":{
  },
  "f:size":{
  },
  "f:volumeSpec":{
    ".":{
    },
    "f:persistentVolumeClaim":{
      ".":{
      },
      "f:resources":{
        ".":{
        },
        "f:requests":{
          ".":{
          },
          "f:storage":{
          }
        }
      }
    }
  }
}

```

```

    }
  }
},
"f:secretsName":{
},
"f:sslInternalSecretName":{
},
"f:sslSecretName":{
},
"f:vaultSecretName":{
}
}
},
{
"manager":"percona-xtradb-cluster-operator",
"operation":"Update",
"apiVersion":"pxc.percona.com/v1",
"time":"2020-05-27T22:32:31Z",
"fieldsType":"FieldsV1",
"fieldsV1":{
"f:spec":{
"f:backup":{
"f:storages":{
"f:fs-pvc":{
"f:podSecurityContext":{
".":{
},
"f:fsGroup":{
},
"f:supplementalGroups":{
}
},
"f:s3":{
".":{
},
"f:bucket":{
},
"f:credentialsSecret":{
}
},
"f:s3-us-west":{
"f:podSecurityContext":{
".":{
},
"f:fsGroup":{
},
"f:supplementalGroups":{
}
}
}
}
},
"f:pmm":{
"f:resources":{
}
},
"f:proxysql":{
"f:podSecurityContext":{
".":{
},
"f:fsGroup":{
},
"f:supplementalGroups":{
}
},
"f:sslInternalSecretName":{
},
"f:sslSecretName":{
},
"f:volumeSpec":{
"f:persistentVolumeClaim":{
"f:accessModes":{
}
}
}
},
"f:pxc":{

```

```

        "f:podSecurityContext":{
          ".":{
            },
            "f:fsGroup":{
            },
            "f:supplementalGroups":{
            }
          },
          "f:sslInternalSecretName":{
          },
          "f:sslSecretName":{
          },
          "f:vaultSecretName":{
          },
          "f:volumeSpec":{
            "f:persistentVolumeClaim":{
              "f:accessModes":{
              }
            }
          }
        },
        "f:status":{
          ".":{
            },
            "f:conditions":{
            },
            "f:host":{
            },
            "f:observedGeneration":{
            },
            "f:proxysql":{
              ".":{
                },
                "f:ready":{
                },
                "f:size":{
                },
                "f:status":{
                }
              },
              "f:pxc":{
                ".":{
                  },
                  "f:ready":{
                  },
                  "f:size":{
                  },
                  "f:status":{
                  }
                },
                "f:state":{
                }
              }
            }
          }
        }
      ]
    ]
  }
}

```

15.3.4 Get status of Percona XtraDB Cluster

Description:

Gets all information about the specified Percona XtraDB Cluster

Kubectl Command:

```
$ kubectl get pxc/cluster1 -o json
```

URL:

```
https://$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters/cluster1
```

Authentication:

```
Authorization: Bearer $KUBE_TOKEN
```

cURL Request:

```
$ curl -k -v -XGET "https://$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters/cluster1" \  
-H "Accept: application/json" \  
-H "Authorization: Bearer $KUBE_TOKEN"
```

Request Body:

```
None
```

Response:

Example

```

{
  "apiVersion": "pxc.percona.com/v1",
  "kind": "PerconaXtraDBCluster",
  "metadata": {
    "annotations": {
      "kubectl.kubernetes.io/last-applied-configuration": "{\"apiVersion\":\"pxc.percona.com/v1\",\"kind\":\"PerconaXtraDBCluster\",\"metadata\":{\"annotations\":{},\"creationTimestamp\":\"2020-05-27T22:23:58Z\"},\"finalizers\":[\"delete-pxc-pods-in-order\"],\"generation\":1,\"managedFields\":[{\"apiVersion\":\"pxc.percona.com/v1-5-0\",\"fieldsType\":\"FieldsV1\",\"fieldsV1\":{\"f:metadata\":{\"f:finalizers\":{\"\"\":{}}},\"f:spec\":{\"f:allowUnsafeConfigurations\":{\"\"\":{}}},\"f:backup\":{\"f:image\":{\"\"\":{}}},\"f:schedule\":{\"\"\":{}}},\"f:serviceAccountName\":{\"\"\":{}}},\"f:storage\":{\"f:fs-pvc\":{\"f:fs-pvc\":{\"f:type\":{\"\"\":{}}},\"f:volume\":{\"f:antiAffinityTopologyKey\":{\"\"\":{}}},\"f:persistentVolumeClaim\":{\"f:accessModes\":{\"\"\":{}}},\"f:resources\":{\"f:requests\":{\"\"\":{}}},\"f:storage\":{\"f:s3-us-west\":{\"f:s3-us-west\":{\"f:s3\":{\"f:bucket\":{\"\"\":{}}},\"f:credentialsSecret\":{\"\"\":{}}},\"f:region\":{\"\"\":{}}},\"f:type\":{\"\"\":{}}},\"f:pmm\":{\"f:image\":{\"\"\":{}}},\"f:serverHost\":{\"\"\":{}}},\"f:serverUser\":{\"\"\":{}}},\"f:proxysql\":{\"f:affinity\":{\"f:antiAffinityTopologyKey\":{\"\"\":{}}},\"f:enabled\":{\"\"\":{}}},\"f:gracePeriod\":{\"\"\":{}}},\"f:podDisruptionBudget\":{\"f:storage\":{\"f:fs-pvc\":{\"f:fs-pvc\":{\"f:volumeSpec\":{\"f:antiAffinityTopologyKey\":{\"\"\":{}}},\"f:gracePeriod\":{\"\"\":{}}},\"f:image\":{\"\"\":{}}},\"f:podDisruptionBudget\":{\"\"\":{}}},\"f:maxUnavailable\":{\"\"\":{}}},\"f:resources\":{\"f:size\":{\"\"\":{}}},\"f:volumeSpec\":{\"f:antiAffinityTopologyKey\":{\"\"\":{}}},\"f:resources\":{\"f:requests\":{\"\"\":{}}},\"f:storage\":{\"f:fs-pvc\":{\"f:fs-pvc\":{\"f:volumeSpec\":{\"f:antiAffinityTopologyKey\":{\"\"\":{}}},\"f:gracePeriod\":{\"\"\":{}}},\"f:image\":{\"\"\":{}}},\"f:podDisruptionBudget\":{\"\"\":{}}},\"f:maxUnavailable\":{\"\"\":{}}},\"f:resources\":{\"f:size\":{\"\"\":{}}},\"f:volumeSpec\":{\"f:antiAffinityTopologyKey\":{\"\"\":{}}},\"f:resources\":{\"f:requests\":{\"\"\":{}}},\"f:storage\":{\"\"\":{}}},\"f:secretsName\":{\"\"\":{}}},\"f:ssInternalSecretName\":{\"\"\":{}}},\"f:ssSecretName\":{\"\"\":{}}},\"f:vaultSecretName\":{\"\"\":{}}},\"manager\":{\"kubectl\",\"operation\":\"Update\",\"time\":\"2020-05-27T22:23:58Z\"},\"apiVersion\":\"pxc.percona.com/v1\",\"fieldsType\":\"FieldsV1\",\"fieldsV1\":{\"f:spec\":{\"f:backup\":{\"f:storages\":{\"f:fs-pvc\":{\"f:fs-pvc\":{\"f:podSecurityContext\":{\"\"\":{}}},\"f:fsGroup\":{\"\"\":{}}},\"f:supplementalGroups\":{\"\"\":{}}},\"f:s3\":{\"f:s3\":{\"f:bucket\":{\"\"\":{}}},\"f:credentialsSecret\":{\"\"\":{}}},\"f:s3-us-west\":{\"f:podSecurityContext\":{\"\"\":{}}},\"f:fsGroup\":{\"\"\":{}}},\"f:supplementalGroups\":{\"\"\":{}}},\"f:ssInternalSecretName\":{\"\"\":{}}},\"f:ssSecretName\":{\"\"\":{}}},\"f:storage\":{\"f:fs-pvc\":{\"f:fs-pvc\":{\"f:volumeSpec\":{\"f:antiAffinityTopologyKey\":{\"\"\":{}}},\"f:gracePeriod\":{\"\"\":{}}},\"f:image\":{\"\"\":{}}},\"f:podDisruptionBudget\":{\"\"\":{}}},\"f:maxUnavailable\":{\"\"\":{}}},\"f:resources\":{\"f:size\":{\"\"\":{}}},\"f:volumeSpec\":{\"f:antiAffinityTopologyKey\":{\"\"\":{}}},\"f:resources\":{\"f:requests\":{\"\"\":{}}},\"f:storage\":{\"\"\":{}}},\"f:secretsName\":{\"\"\":{}}},\"f:ssInternalSecretName\":{\"\"\":{}}},\"f:ssSecretName\":{\"\"\":{}}},\"f:vaultSecretName\":{\"\"\":{}}},\"manager\":{\"kubectl\",\"operation\":\"Update\",\"time\":\"2020-05-27T23:06:47Z\"},\"name\":\"cluster1\",\"namespace\":\"default\",\"resourceVersion\":\"15878\",\"selfLink\":\"/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters/cluster1\",\"uid\":\"e9115e2a-49df-4ebf-9dab-fa5a550208d3\"},\"spec\":{\"allowUnsafeConfigurations\":true,\"backup\":{\"f:image\":{\"\"\":{}}},\"f:storage\":{\"f:s3-us-west\":{\"f:s3-us-west\":{\"f:s3\":{\"f:bucket\":{\"\"\":{}}},\"f:credentialsSecret\":{\"\"\":{}}},\"f:region\":{\"\"\":{}}},\"f:type\":{\"\"\":{}}},\"f:pmm\":{\"f:image\":{\"\"\":{}}},\"f:serverHost\":{\"\"\":{}}},\"f:serverUser\":{\"\"\":{}}},\"f:proxysql\":{\"f:affinity\":{\"f:antiAffinityTopologyKey\":{\"\"\":{}}},\"f:enabled\":true,\"f:gracePeriod\":30,\"f:image\":{\"percona/percona-xtradb-cluster-operator:1.5.0-pmm\"},\"f:podDisruptionBudget\":{\"f:storage\":{\"f:fs-pvc\":{\"f:fs-pvc\":{\"f:volumeSpec\":{\"f:antiAffinityTopologyKey\":{\"\"\":{}}},\"f:gracePeriod\":{\"\"\":{}}},\"f:image\":{\"\"\":{}}},\"f:podDisruptionBudget\":{\"\"\":{}}},\"f:maxUnavailable\":{\"\"\":{}}},\"f:resources\":{\"f:size\":{\"\"\":{}}},\"f:volumeSpec\":{\"f:antiAffinityTopologyKey\":{\"\"\":{}}},\"f:resources\":{\"f:requests\":{\"\"\":{}}},\"f:storage\":{\"\"\":{}}},\"f:secretsName\":{\"\"\":{}}},\"f:ssInternalSecretName\":{\"\"\":{}}},\"f:ssSecretName\":{\"\"\":{}}},\"f:vaultSecretName\":{\"\"\":{}}},\"keyring-secret-vault\":{\"\"\":{}}},\"status\":{\"conditions\":[{\"lastTransitionTime\":\"2020-05-27T22:23:58Z\",\"status\":\"True\",\"type\":\"Initializing\"},{\"lastTransitionTime\":\"2020-05-27T22:25:43Z\",\"status\":\"True\",\"type\":\"Ready\"}],\"host\":\"cluster1-proxysql.default\",\"observedGeneration\":1,\"proxysql\":{\"ready\":3,\"size\":5,\"status\":\"ready\"},\"pxc\":{\"ready\":3,\"size\":5,\"status\":\"ready\"},\"state\":\"ready\"}}]\"}
    },
    "creationTimestamp": "2020-05-27T22:23:58Z",
    "finalizers": [
      "delete-pxc-pods-in-order"
    ],
    "generation": 6,
    "managedFields": [
      {
        "apiVersion": "pxc.percona.com/v1-5-0",
        "fieldsType": "FieldsV1",
        "fieldsV1": {
          "f:metadata": {
            "f:finalizers": {
              "": {
            }
          }
        },
        "f:spec": {
          "": {
        },
        "f:allowUnsafeConfigurations": {
          "": {
        },
        "f:backup": {
          "": {
        },
        "f:schedule": {
          "": {
        },
        "f:serviceAccountName": {
          "": {
        },
        "f:storages": {
          "": {
        },
        "f:fs-pvc": {
          "": {
        },
        "f:type": {
          "": {
        },
        "f:volume": {
          "": {
        },
        "f:persistentVolumeClaim": {
          "": {
        },
        "f:accessModes": {
          "": {
        },
        "f:resources": {
          "": {
        },
        "f:requests": {
          "": {
        }
      }
    }
  }

```



```

        },
        "f:storage":{
            }
        }
    }
},
"f:s3-us-west":{
    ".":{
    },
    "f:s3":{
        ".":{
        },
        "f:bucket":{
        },
        "f:credentialsSecret":{
        },
        "f:region":{
        }
    },
    "f:type":{
    }
}
},
"f:pmm":{
    ".":{
    },
    "f:image":{
    },
    "f:serverHost":{
    },
    "f:serverUser":{
    }
},
"f:proxysql":{
    ".":{
    },
    "f:affinity":{
        ".":{
        },
        "f:antiAffinityTopologyKey":{
        }
    },
    "f:enabled":{
    },
    "f:gracePeriod":{
    },
    "f:image":{
    },
    "f:podDisruptionBudget":{
        ".":{
        },
        "f:maxUnavailable":{
        }
    },
    "f:resources":{
    },
    "f:volumeSpec":{
        ".":{
        },
        "f:persistentVolumeClaim":{
            ".":{
            },
            "f:resources":{
                ".":{
                },
                "f:requests":{
                    ".":{
                    },
                    "f:storage":{
                    }
                }
            }
        }
    }
}
}
}
}

```

```

    }
  }
},
"f:pxc":{
  ".":{"
    },
    "f:affinity":{
      ".":{"
        },
        "f:antiAffinityTopologyKey":{
          }
        },
        "f:gracePeriod":{
          },
          "f:podDisruptionBudget":{
            ".":{"
              },
              "f:maxUnavailable":{
                }
              },
              "f:resources":{
                },
                "f:volumeSpec":{
                  ".":{"
                    },
                    "f:persistentVolumeClaim":{
                      ".":{"
                        },
                        "f:resources":{
                          ".":{"
                            },
                            "f:requests":{
                              ".":{"
                                },
                                "f:storage":{
                                  }
                                }
                              }
                            }
                          }
                        }
                      }
                    }
                  }
                }
              }
            }
          }
        },
        "f:secretsName":{
          },
          "f:ssInternalSecretName":{
            },
            "f:ssSecretName":{
              },
              "f:vaultSecretName":{
                }
              }
            }
          }
        },
        "manager":"kubectl",
        "operation":"Update",
        "time":"2020-05-27T22:23:58Z"
      },
      {
        "apiVersion":"pxc.percona.com/v1",
        "fieldsType":"FieldsV1",
        "fieldsV1":{
          "metadata":{
            "f:annotations":{
              ".":{"
                },
                "f:kubectl.kubernetes.io/last-applied-configuration":{
                  }
                }
              }
            },
            "f:spec":{
              "f:backup":{
                "f:image":{
                  }
                },
                "f:proxysql":{
                  "f:size":{
                    }
                  }
                },
                "f:pxc":{
                  "f:image":{

```

```

    },
    "f:size":{
      }
    }
  },
  "manager":"kubectl",
  "operation":"Update",
  "time":"2020-05-27T23:38:49Z"
},
{
  "apiVersion":"pxc.percona.com/v1",
  "fieldsType":"FieldsV1",
  "fieldsV1":{
    "f:spec":{
      "f:backup":{
        "f:storages":{
          "f:fs-pvc":{
            "f:podSecurityContext":{
              ".":{"
            },
            "f:fsGroup":{
            },
            "f:supplementalGroups":{
            }
          },
          "f:s3":{
            ".":{"
          },
          "f:bucket":{
          },
          "f:credentialsSecret":{
          }
        },
        "f:s3-us-west":{
          "f:podSecurityContext":{
            ".":{"
          },
          "f:fsGroup":{
          },
          "f:supplementalGroups":{
          }
        }
      }
    },
    "f:pmm":{
      "f:resources":{
      }
    },
    "f:proxysql":{
      "f:podSecurityContext":{
        ".":{"
      },
      "f:fsGroup":{
      },
      "f:supplementalGroups":{
      }
    },
    "f:sslInternalSecretName":{
    },
    "f:sslSecretName":{
    },
    "f:volumeSpec":{
      "f:persistentVolumeClaim":{
        "f:accessModes":{
        }
      }
    }
  },
  "f:pxc":{
    "f:podSecurityContext":{
      ".":{"
    },
    "f:fsGroup":{
    },
    "f:supplementalGroups":{
    }
  },
}

```

```

    "f:sslInternalSecretName":{
      },
      "f:sslSecretName":{
        },
        "f:vaultSecretName":{
          },
          "f:volumeSpec":{
            "f:persistentVolumeClaim":{
              "f:accessModes":{
                }
              }
            }
          }
        },
        "f:status":{
          ".":{
            },
            "f:conditions":{
              },
              "f:host":{
                },
                "f:message":{
                  },
                  "f:observedGeneration":{
                    },
                    "f:proxysql":{
                      ".":{
                        },
                        "f:ready":{
                          },
                          "f:size":{
                            },
                            "f:status":{
                              }
                            }
                          },
                          "f:pxc":{
                            ".":{
                              },
                              "f:message":{
                                },
                                "f:ready":{
                                  },
                                  "f:size":{
                                    },
                                    "f:status":{
                                      }
                                    }
                                  },
                                  "f:state":{
                                    }
                                  }
                                }
                              },
                              "manager":"percona-xtradb-cluster-operator",
                              "operation":"Update",
                              "time":"2020-05-28T10:42:00Z"
                            }
                          },
                          "name":"cluster1",
                          "namespace":"default",
                          "resourceVersion":"35660",
                          "selfLink":"/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters/cluster1",
                          "uid":"e9115e2a-49df-4ebf-9dab-fa5a550208d3"
                        },
                        "spec":{
                          "allowUnsafeConfigurations":true,
                          "backup":{
                            "image":"percona/percona-xtradb-cluster-operator:1.5.0-pxc8.0-debug-backup",
                            "schedule":[
                              {
                                "keep":3,
                                "name":"sat-night-backup",
                                "schedule":"0 0 * * 6",
                                "storageName":"s3-us-west"
                              },
                              {
                                "keep":5,
                                "name":"daily-backup",
                                "schedule":"0 0 * * *",
                                "storageName":"fs-pvc"
                              }
                            ]
                          },
                          "serviceAccountName":"percona-xtradb-cluster-operator",

```

```

"storages":{
  "fs-pvc":{
    "type":"fiLesystem",
    "volume":{
      "persistentVolumeClaim":{
        "accessModes":[
          "ReadWriteOnce"
        ],
        "resources":{
          "requests":{
            "storage":"6Gi"
          }
        }
      }
    }
  },
  "s3-us-west":{
    "s3":{
      "bucket":"S3-BACKUP-BUCKET-NAME-HERE",
      "credentialsSecret":"my-cluster-name-backup-s3",
      "region":"us-west-2"
    },
    "type":"s3"
  }
},
"pmm":{
  "enabled":false,
  "image":"percona/percona-xtradb-cluster-operator:1.5.0-pmm",
  "serverHost":"monitoring-service",
  "serverUser":"pmm"
},
"proxysql":{
  "affinity":{
    "antiAffinityTopologyKey":"none"
  },
  "enabled":true,
  "gracePeriod":30,
  "image":"percona/percona-xtradb-cluster-operator:1.5.0-proxysql",
  "podDisruptionBudget":{
    "maxUnavailable":1
  },
  "resources":{
  },
  "size":3,
  "volumeSpec":{
    "persistentVolumeClaim":{
      "resources":{
        "requests":{
          "storage":"2Gi"
        }
      }
    }
  }
},
"pxc":{
  "affinity":{
    "antiAffinityTopologyKey":"none"
  },
  "gracePeriod":600,
  "image":"percona/percona-xtradb-cluster-operator:1.5.0-pxc8.0-debug",
  "podDisruptionBudget":{
    "maxUnavailable":1
  },
  "resources":{
  },
  "size":3,
  "volumeSpec":{
    "persistentVolumeClaim":{
      "resources":{
        "requests":{
          "storage":"6Gi"
        }
      }
    }
  }
},
"secretsName":"my-cluster-secrets",
"sslInternalSecretName":"my-cluster-ssl-internal",
"sslSecretName":"my-cluster-ssl",
"vaultSecretName":"keyring-secret-vault"
},
"status":{
  "conditions":[
    {
      "lastTransitionTime":"2020-05-27T22:25:43Z",
      "status":"True",
      "type":"Ready"
    },
    {
      "lastTransitionTime":"2020-05-27T23:06:48Z",
      "status":"True",
      "type":"Initializing"
    },
    {
      "lastTransitionTime":"2020-05-27T23:08:58Z",
      "message":"ProxySQL upgrade error: context deadline exceeded",
      "reason":"ErrorReconcile",

```

```

    "status":"True",
    "type":"Error"
  },
  {
    "lastTransitionTime":"2020-05-27T23:08:59Z",
    "status":"True",
    "type":"Initializing"
  },
  {
    "lastTransitionTime":"2020-05-27T23:29:59Z",
    "status":"True",
    "type":"Ready"
  },
  {
    "lastTransitionTime":"2020-05-27T23:30:04Z",
    "status":"True",
    "type":"Initializing"
  },
  {
    "lastTransitionTime":"2020-05-27T23:35:27Z",
    "status":"True",
    "type":"Ready"
  },
  {
    "lastTransitionTime":"2020-05-27T23:35:42Z",
    "status":"True",
    "type":"Initializing"
  },
  {
    "lastTransitionTime":"2020-05-27T23:47:00Z",
    "status":"True",
    "type":"Ready"
  },
  {
    "lastTransitionTime":"2020-05-27T23:47:05Z",
    "status":"True",
    "type":"Initializing"
  },
  {
    "lastTransitionTime":"2020-05-28T09:58:25Z",
    "status":"True",
    "type":"Ready"
  },
  {
    "lastTransitionTime":"2020-05-28T09:58:31Z",
    "status":"True",
    "type":"Initializing"
  },
  {
    "lastTransitionTime":"2020-05-28T10:03:54Z",
    "status":"True",
    "type":"Ready"
  },
  {
    "lastTransitionTime":"2020-05-28T10:04:14Z",
    "status":"True",
    "type":"Initializing"
  },
  {
    "lastTransitionTime":"2020-05-28T10:15:28Z",
    "status":"True",
    "type":"Ready"
  },
  {
    "lastTransitionTime":"2020-05-28T10:15:38Z",
    "status":"True",
    "type":"Initializing"
  },
  {
    "lastTransitionTime":"2020-05-28T10:26:56Z",
    "status":"True",
    "type":"Ready"
  },
  {
    "lastTransitionTime":"2020-05-28T10:27:01Z",
    "status":"True",
    "type":"Initializing"
  },
  {
    "lastTransitionTime":"2020-05-28T10:38:28Z",
    "status":"True",
    "type":"Ready"
  },
  {
    "lastTransitionTime":"2020-05-28T10:38:33Z",
    "status":"True",
    "type":"Initializing"
  }
],
"host":"cluster1-proxysql.default",
"message": [
  "PXC: pxc: back-off 5m0s restarting failed container=pxc pod=cluster1-pxc-1_default(5b9b16e6-d0f8-4c97-a2d0-294feb9d014b); pxc: back-off 5m0s restarting failed container=pxc pod=cluster1-pxc-2_default(b8ebdd7-42f0-440b-aa5e-509d28926a5e); pxc: back-off 5m0s restarting failed container=pxc pod=cluster1-pxc-4_default(2dce12f2-9ebc-419c-a92a-9cec68912004); "
],
"observedGeneration":6,
"proxysql":{
  "ready":3,
  "size":3,
  "status":"ready"
}
},

```

```

    "pxc":{
      "message":"pxc: back-off 5m0s restarting failed container=pxc pod=cluster1-pxc-1_default(5b9b16e6-d0f8-4c97-a2d0-294feb9d014b); pxc: back-off 5m0s restarting failed container=pxc
pod=cluster1-pxc-2_default(b8ebdd7-42f0-440b-aa5e-509d28926a5e); pxc: back-off 5m0s restarting failed container=pxc pod=cluster1-pxc-4_default(2dce12f2-9ebc-419c-a92a-9cec68912004); ",
      "ready":2,
      "size":3,
      "status":"initializing"
    },
    "state":"initializing"
  }
}

```

15.3.5 Scale up/down Percona XtraDB Cluster

Description:

Increase or decrease the size of the Percona XtraDB Cluster nodes to fit the current high availability needs

Kubect1 Command:

```

$ kubectl patch pxc cluster1 --type=merge --patch '{
"spec": {"pxc":{"size": "5" }
}}'

```

URL:

[https://\\$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters/cluster1](https://$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters/cluster1)

Authentication:

Authorization: Bearer \$KUBE_TOKEN

cURL Request:

```

$ curl -k -v -XPATCH "https://$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters/cluster1" \
-H "Authorization: Bearer $KUBE_TOKEN" \
-H "Content-Type: application/merge-patch+json" \
-H "Accept: application/json" \
-d '{
  "spec": {"pxc":{"size": "5" }
}'

```

Request Body:

Example

```

{
  "spec": {"pxc":{"size": "5" }
}}

```


Input:

spec:

pxc

1. size (Int or String, Defaults: 3): Specify the size of the Percona XtraDB Cluster to scale up or down to

Response:

 **Example**

```

{
  "apiVersion": "pxc.percona.com/v1",
  "kind": "PerconaXtraDBCluster",
  "metadata": {
    "annotations": {
      "kubectrl.kubernetes.io/last-applied-configuration": "{\n  \"apiVersion\": \"pxc.percona.com/v1-5-0\", \"kind\": \"PerconaXtraDBCluster\", \"metadata\": {\n    \"annotations\": {\n      \"finalizers\": [\n        [\n          {\n            \"delete-pxc-pods-in-order\": {\n              \"name\": \"cluster1\", \"namespace\": \"default\", \"spec\": {\n                \"allowUnsafeConfigurations\": true, \"backup\": {\n                  \"image\": \"percona/percona-xtradb-cluster-operator:1.5.0-pxc8.0-backup\", \"schedule\": {\n                    \"keep\": 3, \"name\": \"sat-night-backup\", \"schedule\": \"0 0 * * 6\", \"storageName\": \"s3-us-west\", \"keep\": 5, \"name\": \"daily-backup\", \"schedule\": \"0 0 * * *\", \"storageName\": \"fs-pvc\"}, \"serviceAccountName\": \"percona-xtradb-cluster-operator\", \"storages\": {\n                      \"fs-pvc\": {\n                        \"type\": \"filesystem\", \"volume\": {\n                          \"persistentVolumeClaim\": {\n                            \"accessModes\": [\"ReadWriteOnce\"], \"resources\": {\n                              \"requests\": {\n                                \"storage\": \"6Gi\"}}}}}, \"s3-us-west\": {\n                      \"s3\": {\n                        \"bucket\": \"S3-BACKUP-BUCKET-NAME-HERE\", \"credentialsSecret\": \"my-cluster-name-backup-s3\", \"region\": \"us-west-2\"}, \"type\": \"s3\"}}}, \"pmm\": {\n                      \"enabled\": false, \"image\": \"percona/percona-xtradb-cluster-operator:1.5.0-pmm\", \"serverHost\": \"monitoring-service\", \"serverUser\": \"pmm\"}, \"proxysql\": {\n                      \"affinity\": {\n                        \"antiAffinityTopologyKey\": \"none\"}, \"enabled\": true, \"gracePeriod\": 30, \"image\": \"percona/percona-xtradb-cluster-operator:1.5.0-proxysql\", \"podDisruptionBudget\": {\n                        \"maxUnavailable\": 1}, \"resources\": {\n                          \"requests\": null, \"size\": 3, \"volumeSpec\": {\n                            \"persistentVolumeClaim\": {\n                              \"resources\": {\n                                \"requests\": {\n                                  \"storage\": \"2Gi\"}}}}}}}, \"pxc\": {\n                      \"affinity\": {\n                        \"antiAffinityTopologyKey\": \"none\"}, \"gracePeriod\": 600, \"image\": \"percona/percona-xtradb-cluster:8.0.19-10.1\", \"podDisruptionBudget\": {\n                        \"maxUnavailable\": 1}, \"resources\": {\n                          \"requests\": null, \"size\": 3, \"volumeSpec\": {\n                            \"persistentVolumeClaim\": {\n                              \"resources\": {\n                                \"requests\": {\n                                  \"storage\": \"6Gi\"}}}}}}}, \"secretsName\": \"my-cluster-secrets\", \"sslInternalSecretName\": \"my-cluster-ssl-internal\", \"sslSecretName\": \"my-cluster-ssl\", \"updateStrategy\": \"RollingUpdate\", \"vaultSecretName\": \"keyring-secret-vault\"}}\n}"}",
      "creationTimestamp": "2020-06-01T16:50:05Z",
      "finalizers": [
        "delete-pxc-pods-in-order"
      ]
    },
    "generation": 4,
    "managedFields": [
      {
        "apiVersion": "pxc.percona.com/v1-5-0",
        "fieldsType": "FieldsV1",
        "fieldsV1": {
          "f:metadata": {
            "f:annotations": {
              ".": {
                "kubectrl.kubernetes.io/last-applied-configuration": {
                }
              }
            },
            "f:finalizers": {
              ".": {
                "delete-pxc-pods-in-order": {
                }
              }
            }
          },
          "f:spec": {
            ".": {
              "f:allowUnsafeConfigurations": {
                "f:backup": {
                  "f:image": {
                  },
                  "f:schedule": {
                  },
                  "f:serviceAccountName": {
                  },
                  "f:storages": {
                    "f:fs-pvc": {
                      "f:type": {
                      },
                      "f:volume": {
                      },
                      "f:persistentVolumeClaim": {
                      },
                      "f:accessModes": {
                      },
                      "f:resources": {
                      },
                      "f:requests": {
                      },
                      "f:storage": {
                      }
                    }
                  }
                }
              }
            }
          }
        }
      ]
    }
  }
}

```



```

    }
  },
  "f:s3-us-west":{
    ".":{
      },
      "f:s3":{
        ".":{
          },
          "f:bucket":{
            },
            "f:credentialsSecret":{
              },
              "f:region":{
                }
              },
              "f:type":{
                }
            }
          }
        },
        "f:pmm":{
          ".":{
            },
            "f:image":{
              },
              "f:serverHost":{
                },
                "f:serverUser":{
                  }
                }
              },
              "f:proxysql":{
                ".":{
                  },
                  "f:affinity":{
                    ".":{
                      },
                      "f:antiAffinityTopologyKey":{
                        }
                      },
                      "f:enabled":{
                        },
                        "f:gracePeriod":{
                          },
                          "f:image":{
                            },
                            "f:podDisruptionBudget":{
                              ".":{
                                },
                                "f:maxUnavailable":{
                                  }
                                },
                                "f:resources":{
                                  },
                                  "f:size":{
                                    },
                                    "f:volumeSpec":{
                                      ".":{
                                        },
                                        "f:persistentVolumeClaim":{
                                          ".":{
                                            },
                                            "f:resources":{
                                              ".":{
                                                },
                                                "f:requests":{
                                                  ".":{
                                                    },
                                                    "f:storage":{
                                                      }
                                                    }
                                                  }
                                                }
                                              }
                                            }
                                          }
                                        }
                                      },
                                      "f:pxc":{

```

```

      ".":{"
        },
        "f:affinity":{"
          ".":{"
            },
            "f:antiAffinityTopologyKey":{"
              }
            },
            "f:gracePeriod":{"
              },
              "f:podDisruptionBudget":{"
                ".":{"
                  },
                  "f:maxUnavailable":{"
                    }
                  },
                  "f:resources":{"
                    },
                    "f:volumeSpec":{"
                      ".":{"
                        },
                        "f:persistentVolumeClaim":{"
                          ".":{"
                            },
                            "f:resources":{"
                              ".":{"
                                },
                                "f:requests":{"
                                  ".":{"
                                    },
                                    "f:storage":{"
                                      }
                                    }
                                  }
                                }
                              }
                            }
                          }
                        }
                      }
                    }
                  },
                  "f:secretsName":{"
                    },
                    "f:sslInternalSecretName":{"
                    },
                    "f:sslSecretName":{"
                    },
                    "f:updateStrategy":{"
                    },
                    "f:vaultSecretName":{"
                    }
                  }
                }
              },
              "manager":"kubectl",
              "operation":"Update",
              "time":"2020-06-01T16:52:30Z"
            },
            {
              "apiVersion":"pxc.percona.com/v1",
              "fieldsType":"FieldsV1",
              "fieldsV1":{"
                "f:spec":{"
                  "f:backup":{"
                    "f:storages":{"
                      "f:fs-pvc":{"
                        "f:podSecurityContext":{"
                          ".":{"
                            },
                            "f:fsGroup":{"
                            },
                            "f:supplementalGroups":{"
                            }
                          },
                          "f:s3":{"
                            ".":{"
                            },
                            "f:bucket":{"
                            },
                            "f:credentialsSecret":{"
                            }
                          }
                        }
                      }
                    }
                  }
                }
              }
            }
          }
        }
      }
    }
  }
}

```

```

    },
    "f:s3-us-west":{
      "f:podSecurityContext":{
        ".":{
          },
          "f:fsGroup":{
            },
            "f:supplementalGroups":{
              }
            }
          }
        },
        "f:pmm":{
          "f:resources":{
            }
          },
          "f:proxysql":{
            "f:podSecurityContext":{
              ".":{
                },
                "f:fsGroup":{
                  },
                  "f:supplementalGroups":{
                    }
                  }
                },
                "f:sslInternalSecretName":{
                  },
                  "f:sslSecretName":{
                    },
                    "f:volumeSpec":{
                      "f:persistentVolumeClaim":{
                        "f:accessModes":{
                          }
                        }
                      }
                    },
                    "f:pxc":{
                      "f:podSecurityContext":{
                        ".":{
                          },
                          "f:fsGroup":{
                            },
                            "f:supplementalGroups":{
                              }
                            }
                          },
                          "f:sslInternalSecretName":{
                            },
                            "f:sslSecretName":{
                              },
                              "f:vaultSecretName":{
                                },
                                "f:volumeSpec":{
                                  "f:persistentVolumeClaim":{
                                    "f:accessModes":{
                                      }
                                    }
                                  }
                                }
                              },
                              "f:status":{
                                ".":{
                                  },
                                  "f:conditions":{
                                    },
                                    "f:host":{
                                      },
                                      "f:observedGeneration":{
                                        },
                                        "f:proxysql":{
                                          ".":{
                                            },
                                            "f:ready":{
                                              },
                                              "f:size":{
                                                },
                                                },

```

```

        "f:status":{
            }
        },
        "f:pxc":{
            ".":{
            },
            "f:ready":{
            },
            "f:size":{
            },
            "f:status":{
            }
        },
        "f:state":{
            }
        }
    },
    "manager": "percona-xtradb-cluster-operator",
    "operation": "Update",
    "time": "2020-06-03T15:32:11Z"
},
{
    "apiVersion": "pxc.percona.com/v1",
    "fieldsType": "FieldsV1",
    "fieldsV1": {
        "f:spec": {
            "f:pxc": {
                "f:image": {
                },
                "f:size": {
                }
            }
        }
    },
    "manager": "kubectl",
    "operation": "Update",
    "time": "2020-06-03T15:32:14Z"
}
],
"name": "cluster1",
"namespace": "default",
"resourceVersion": "129605",
"selfLink": "/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters/cluster1",
"uid": "15e5e7d6-10b2-46cf-85d0-d3fdea3412ca"
},
"spec": {
    "allowUnsafeConfigurations": true,
    "backup": {
        "image": "percona/percona-xtradb-cluster-operator:1.5.0-pxc8.0-backup",
        "schedule": [
            {
                "keep": 3,
                "name": "sat-night-backup",
                "schedule": "0 0 * * 6",
                "storageName": "s3-us-west"
            },
            {
                "keep": 5,
                "name": "daily-backup",
                "schedule": "0 0 * * *",
                "storageName": "fs-pvc"
            }
        ]
    },
    "serviceAccountName": "percona-xtradb-cluster-operator",
    "storages": {
        "fs-pvc": {
            "type": "filesystem",
            "volume": {
                "persistentVolumeClaim": {
                    "accessModes": [
                        "ReadWriteOnce"
                    ],
                    "resources": {
                        "requests": {
                            "storage": "6Gi"
                        }
                    }
                }
            }
        },
        "s3-us-west": {
            "s3": {
                "bucket": "S3-BACKUP-BUCKET-NAME-HERE",
                "credentialsSecret": "my-cluster-name-backup-s3",
                "region": "us-west-2"
            },
            "type": "s3"
        }
    }
},
"pmm": {
    "enabled": false,

```

```

    "image": "percona/percona-xtradb-cluster-operator:1.5.0-pmm",
    "serverHost": "monitoring-service",
    "serverUser": "pmm"
  },
  "proxysql": {
    "affinity": {
      "antiAffinityTopologyKey": "none"
    },
    "enabled": true,
    "gracePeriod": 30,
    "image": "percona/percona-xtradb-cluster-operator:1.5.0-proxysql",
    "podDisruptionBudget": {
      "maxUnavailable": 1
    },
    "resources": {
      "requests": null
    },
    "size": 3,
    "volumeSpec": {
      "persistentVolumeClaim": {
        "resources": {
          "requests": {
            "storage": "2Gi"
          }
        }
      }
    }
  },
  "pxc": {
    "affinity": {
      "antiAffinityTopologyKey": "none"
    },
    "gracePeriod": 600,
    "image": "percona/percona-xtradb-cluster:5.7.30-31.43",
    "podDisruptionBudget": {
      "maxUnavailable": 1
    },
    "resources": {
      "requests": null
    },
    "size": 5,
    "volumeSpec": {
      "persistentVolumeClaim": {
        "resources": {
          "requests": {
            "storage": "6Gi"
          }
        }
      }
    }
  },
  "secretsName": "my-cluster-secrets",
  "sslInternalSecretName": "my-cluster-ssl-internal",
  "sslSecretName": "my-cluster-ssl",
  "updateStrategy": "RollingUpdate",
  "vaultSecretName": "keyring-secret-vault"
},
"status": {
  "conditions": [
    {
      "lastTransitionTime": "2020-06-01T16:50:37Z",
      "message": "create newStatefulSetNode: StatefulSet.apps \"cluster1-pxc\" is invalid: spec.updateStrategy: Invalid value: apps.StatefulSetUpdateStrategy{Type:\\\"SmartUpdate\\\", RollingUpdate:(*apps.RollingUpdateStatefulSetStrategy)(nil)}: must be 'RollingUpdate' or 'OnDelete'",
      "reason": "ErrorReconcile",
      "status": "True",
      "type": "Error"
    },
    {
      "lastTransitionTime": "2020-06-01T16:52:31Z",
      "status": "True",
      "type": "Initializing"
    },
    {
      "lastTransitionTime": "2020-06-01T16:55:59Z",
      "status": "True",
      "type": "Ready"
    },
    {
      "lastTransitionTime": "2020-06-01T17:19:15Z",
      "status": "True",
      "type": "Initializing"
    }
  ],
  "host": "cluster1-proxysql.default",
  "observedGeneration": 3,
  "proxysql": {
    "ready": 3,
    "size": 3,
    "status": "ready"
  },
  "pxc": {
    "ready": 1,
    "size": 3,
    "status": "initializing"
  },
  "state": "initializing"
}
}

```

15.3.6 Update Percona XtraDB Cluster image

Description:

Change the image of Percona XtraDB Cluster containers inside the cluster

Kubectl Command:

```
$ kubectl patch pxc cluster1 --type=merge --patch '{
"spec": {"pxc":{ "image": "percona/percona-xtradb-cluster:5.7.30-31.43" }
}}'
```

URL:

https://\$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters/cluster1

Authentication:

Authorization: Bearer \$KUBE_TOKEN

cURL Request:

```
$ curl -k -v -XPATCH "https://$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters/cluster1" \
-H "Authorization: Bearer $KUBE_TOKEN" \
-H "Accept: application/json" \
-H "Content-Type: application/merge-patch+json"
-d '{
"spec": {"pxc":{ "image": "percona/percona-xtradb-cluster:5.7.30-31.43" }
}}'
```

Request Body:

Example

```
{
"spec": {"pxc":{ "image": "percona/percona-xtradb-cluster:5.7.30-31.43" }
}}
```

Input:

spec:

pxc:

1. image (String, min-length: 1) : name of the image to update for Percona XtraDB Cluster

Response:

Example

```

{
  "apiVersion": "pxc.percona.com/v1",
  "kind": "PerconaXtraDBCluster",
  "metadata": {
    "annotations": {
      "kubectrl.kubernetes.io/last-applied-configuration": "{\"apiVersion\":\"pxc.percona.com/v1-5-0\", \"kind\":\"PerconaXtraDBCluster\", \"metadata\": {\"annotations\": {\"finalizers\": [\"delete-pxc-pods-in-order\"], \"name\": \"cluster1\", \"namespace\": \"default\"}, \"spec\": {\"allowUnsafeConfigurations\": true, \"backup\": {\"image\": \"percona/percona-xtradb-cluster-operator:1.5.0-pxc8.0-backup\", \"schedule\": {\"keep\": 3, \"name\": \"sat-night-backup\", \"schedule\": \"0 0 * * 6\", \"storageName\": \"s3-us-west\"}, {\"keep\": 5, \"name\": \"daily-backup\", \"schedule\": \"0 0 * * *\", \"storageName\": \"fs-pvc\"}}, \"serviceAccountName\": \"percona-xtradb-cluster-operator\", \"storages\": {\"fs-pvc\": {\"type\": \"filesystem\", \"volume\": {\"persistentVolumeClaim\": {\"accessModes\": [\"ReadWriteOnce\"], \"resources\": {\"requests\": {\"storage\": \"6Gi\"}}}}, \"s3-us-west\": {\"s3\": {\"bucket\": \"S3-BACKUP-BUCKET-NAME-HERE\", \"credentialsSecret\": \"my-cluster-name-backup-s3\", \"region\": \"us-west-2\"}, \"type\": \"s3\"}}, \"pmm\": {\"enabled\": false, \"image\": \"percona/percona-xtradb-cluster-operator:1.5.0-pmm\", \"serverHost\": \"monitoring-service\", \"serverUser\": \"pmm\"}, \"proxysql\": {\"affinity\": {\"antiAffinityTopologyKey\": \"none\"}, \"enabled\": true, \"gracePeriod\": 30, \"image\": \"percona/percona-xtradb-cluster-operator:1.5.0-proxysql\", \"podDisruptionBudget\": {\"maxUnavailable\": 1}, \"resources\": {\"requests\": null, \"size\": 3, \"volumeSpec\": {\"persistentVolumeClaim\": {\"resources\": {\"requests\": {\"storage\": \"2Gi\"}}}}, \"pxc\": {\"affinity\": {\"antiAffinityTopologyKey\": \"none\"}, \"gracePeriod\": 600, \"image\": \"percona/percona-xtradb-cluster:8.0.19-10.1\", \"podDisruptionBudget\": {\"maxUnavailable\": 1}, \"resources\": {\"requests\": null, \"size\": 3, \"volumeSpec\": {\"persistentVolumeClaim\": {\"resources\": {\"requests\": {\"storage\": \"6Gi\"}}}}, \"secretsName\": \"my-cluster-secrets\", \"sslInternalSecretName\": \"my-cluster-ssl-internal\", \"sslSecretName\": \"my-cluster-ssl\", \"updateStrategy\": \"RollingUpdate\", \"vaultSecretName\": \"keyring-secret-vault\"}}}\n"
    },
    "creationTimestamp": "2020-06-01T16:50:05Z",
    "finalizers": [
      "delete-pxc-pods-in-order"
    ],
    "generation": 3,
    "managedFields": [
      {
        "apiVersion": "pxc.percona.com/v1-5-0",
        "fieldsType": "FieldsV1",
        "fieldsV1": {
          "f:metadata": {
            "f:annotations": {
              ".": {
                "kubectrl.kubernetes.io/last-applied-configuration": {
                }
              }
            },
            "f:finalizers": {
              ".": {
                "delete-pxc-pods-in-order": {
                }
              }
            },
            "f:spec": {
              ".": {
                "f:allowUnsafeConfigurations": {
                }
              },
              "f:backup": {
                ".": {
                  "f:image": {
                  },
                  "f:schedule": {
                  },
                  "f:serviceAccountName": {
                  },
                  "f:storages": {
                    ".": {
                      "f:fs-pvc": {
                        ".": {
                          "f:type": {
                          },
                          "f:volume": {
                          },
                          "f:persistentVolumeClaim": {
                            ".": {
                              "f:accessModes": {
                              },
                              "f:resources": {
                                ".": {
                                  "f:requests": {
                                    ".": {
                                      "f:storage": {
                                      }
                                    }
                                  }
                                }
                              }
                            }
                          }
                        }
                      }
                    }
                  }
                }
              }
            }
          }
        }
      }
    ]
  }
}

```

```

    }
  },
  "f:s3-us-west":{
    ".":{
      },
    },
    "f:s3":{
      ".":{
        },
      },
      "f:bucket":{
        },
      },
      "f:credentialsSecret":{
        },
      },
      "f:region":{
        }
      },
    },
    "f:type":{
      }
    }
  },
  },
  "f:pmm":{
    ".":{
      },
    },
    "f:image":{
      },
    },
    "f:serverHost":{
      },
    },
    "f:serverUser":{
      }
    }
  },
  "f:proxysql":{
    ".":{
      },
    },
    "f:affinity":{
      ".":{
        },
      },
      "f:antiAffinityTopologyKey":{
        }
      },
    },
    "f:enabled":{
      },
    },
    "f:gracePeriod":{
      },
    },
    "f:image":{
      },
    },
    "f:podDisruptionBudget":{
      ".":{
        },
      },
      "f:maxUnavailable":{
        }
      },
    },
    "f:resources":{
      },
    },
    "f:size":{
      },
    },
    "f:volumeSpec":{
      ".":{
        },
      },
      "f:persistentVolumeClaim":{
        ".":{
          },
        },
        "f:resources":{
          ".":{
            },
          },
          "f:requests":{
            ".":{
              },
            },
            "f:storage":{
              }
            }
          }
        }
      },
    },
  },
  "f:pxc":{

```



```

      ".":{
      },
      "f:affinity":{
        ".":{
          },
          "f:antiAffinityTopologyKey":{
            }
          },
          "f:gracePeriod":{
            },
            "f:podDisruptionBudget":{
              ".":{
                },
                "f:maxUnavailable":{
                  }
                },
                "f:resources":{
                  },
                  "f:size":{
                    },
                    "f:volumeSpec":{
                      ".":{
                        },
                        "f:persistentVolumeClaim":{
                          ".":{
                            },
                            "f:resources":{
                              ".":{
                                },
                                "f:requests":{
                                  ".":{
                                    },
                                    "f:storage":{
                                      }
                                    }
                                  }
                                }
                              }
                            }
                          },
                          "f:secretsName":{
                            },
                            "f:ssInternalSecretName":{
                              },
                              "f:ssSecretName":{
                                },
                                "f:updateStrategy":{
                                  },
                                  "f:vaultSecretName":{
                                    }
                                  }
                                }
                              },
                              "manager":"kubectl",
                              "operation":"Update",
                              "time":"2020-06-01T16:52:30Z"
                            },
                            {
                              "apiVersion":"pxc.percona.com/v1",
                              "fieldsType":"FieldsV1",
                              "fieldsV1":{
                                "f:spec":{
                                  "f:pxc":{
                                    "f:image":{
                                      }
                                    }
                                  }
                                }
                              },
                              "manager":"kubectl",
                              "operation":"Update",
                              "time":"2020-06-01T17:18:58Z"
                            },
                            {
                              "apiVersion":"pxc.percona.com/v1",
                              "fieldsType":"FieldsV1",
                              "fieldsV1":{
                                "f:spec":{
                                  "f:backup":{
                                    "f:storages":{
                                      "f:fs-pvc":{
                                        "f:podSecurityContext":{
                                          ".":{

```

```

        },
        "f:fsGroup":{
        },
        "f:supplementalGroups":{
        }
    },
    "f:s3":{
        ".":{"
        },
        "f:bucket":{
        },
        "f:credentialsSecret":{
        }
    },
    },
    "f:s3-us-west":{
        "f:podSecurityContext":{
            ".":{"
            },
            "f:fsGroup":{
            },
            "f:supplementalGroups":{
            }
        }
    },
    },
    "f:pmm":{
        "f:resources":{
        }
    },
    },
    "f:proxysql":{
        "f:podSecurityContext":{
            ".":{"
            },
            "f:fsGroup":{
            },
            "f:supplementalGroups":{
            }
        },
        "f:sslInternalSecretName":{
        },
        "f:sslSecretName":{
        },
        "f:volumeSpec":{
            "f:persistentVolumeClaim":{
                "f:accessModes":{
                }
            }
        }
    },
    },
    "f:pxc":{
        "f:podSecurityContext":{
            ".":{"
            },
            "f:fsGroup":{
            },
            "f:supplementalGroups":{
            }
        },
        "f:sslInternalSecretName":{
        },
        "f:sslSecretName":{
        },
        "f:vaultSecretName":{
        },
        "f:volumeSpec":{
            "f:persistentVolumeClaim":{
                "f:accessModes":{
                }
            }
        }
    },
    },
    "f:status":{
        ".":{"
    },
    },

```

```

        "f:conditions":{
        },
        "f:host":{
        },
        "f:message":{
        },
        "f:observedGeneration":{
        },
        "f:proxysql":{
            ".":{
            },
            "f:ready":{
            },
            "f:size":{
            },
            "f:status":{
            }
        },
        "f:pxc":{
            ".":{
            },
            "f:message":{
            },
            "f:ready":{
            },
            "f:size":{
            },
            "f:status":{
            }
        },
        "f:state":{
        }
    },
    "manager":"percona-xtradb-cluster-operator",
    "operation":"Update",
    "time":"2020-06-01T17:21:36Z"
}
],
"name":"cluster1",
"namespace":"default",
"resourceVersion":"41149",
"selfLink":"/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusters/cluster1",
"uid":"15e5e7d6-10b2-46cf-85d0-d3fdea3412ca"
},
"spec":{
    "allowUnsafeConfigurations":true,
    "backup":{
        "image":"percona/percona-xtradb-cluster-operator:1.5.0-pxc8.0-backup",
        "schedule":[
            {
                "keep":3,
                "name":"sat-night-backup",
                "schedule":"0 0 * * 6",
                "storageName":"s3-us-west"
            },
            {
                "keep":5,
                "name":"daily-backup",
                "schedule":"0 0 * * *",
                "storageName":"fs-pvc"
            }
        ]
    },
    "serviceAccountName":"percona-xtradb-cluster-operator",
    "storages":{
        "fs-pvc":{
            "type":"FileSystem",
            "volume":{
                "persistentVolumeClaim":{
                    "accessModes":[
                        "ReadWriteOnce"
                    ],
                    "resources":{
                        "requests":{
                            "storage":"6Gi"
                        }
                    }
                }
            }
        }
    }
},
"s3-us-west":{
    "s3":{
        "bucket":"S3-BACKUP-BUCKET-NAME-HERE",
        "credentialsSecret":"my-cluster-name-backup-s3",
        "region":"us-west-2"
    }
},

```

```

    "type": "s3"
  }
},
"pmm": {
  "enabled": false,
  "image": "percona/percona-xtradb-cluster-operator:1.5.0-pmm",
  "serverHost": "monitoring-service",
  "serverUser": "pmm"
},
"proxysql": {
  "affinity": {
    "antiAffinityTopologyKey": "none"
  },
  "enabled": true,
  "gracePeriod": 30,
  "image": "percona/percona-xtradb-cluster-operator:1.5.0-proxysql",
  "podDisruptionBudget": {
    "maxUnavailable": 1
  },
  "resources": {
    "requests": null
  },
  "size": 3,
  "volumeSpec": {
    "persistentVolumeClaim": {
      "resources": {
        "requests": {
          "storage": "2Gi"
        }
      }
    }
  }
},
"pxc": {
  "affinity": {
    "antiAffinityTopologyKey": "none"
  },
  "gracePeriod": 600,
  "image": "percona/percona-xtradb-cluster:5.7.30-31.43",
  "podDisruptionBudget": {
    "maxUnavailable": 1
  },
  "resources": {
    "requests": null
  },
  "size": 3,
  "volumeSpec": {
    "persistentVolumeClaim": {
      "resources": {
        "requests": {
          "storage": "6Gi"
        }
      }
    }
  }
},
"secretsName": "my-cluster-secrets",
"sslInternalSecretName": "my-cluster-ssl-internal",
"sslSecretName": "my-cluster-ssl",
"updateStrategy": "RollingUpdate",
"vaultSecretName": "keyring-secret-vault"
},
"status": {
  "conditions": [
    {
      "lastTransitionTime": "2020-06-01T16:50:37Z",
      "message": "create newStatefulSetNode: StatefulSet.apps \"cluster1-pxc\" is invalid: spec.updateStrategy: Invalid value: apps.StatefulSetUpdateStrategy{Type: \"SmartUpdate\", RollingUpdate: (*apps.RollingUpdateStatefulSetStrategy)(nil)}: must be 'RollingUpdate' or 'OnDelete'",
      "reason": "ErrorReconcile",
      "status": "True",
      "type": "Error"
    },
    {
      "lastTransitionTime": "2020-06-01T16:52:31Z",
      "status": "True",
      "type": "Initializing"
    },
    {
      "lastTransitionTime": "2020-06-01T16:55:59Z",
      "status": "True",
      "type": "Ready"
    },
    {
      "lastTransitionTime": "2020-06-01T17:19:15Z",
      "status": "True",
      "type": "Initializing"
    }
  ],
  "host": "cluster1-proxysql.default",
  "message": [
    "PXC: pxc: back-off 40s restarting failed container=pxc pod=cluster1-pxc-2_default(87cdf1a8-0fb3-4bc0-b50d-f66a0a73c087): "
  ],
  "observedGeneration": 3,
  "proxysql": {
    "ready": 3,
    "size": 3,
    "status": "ready"
  },
  "pxc": {

```

```

    "message":"pxc: back-off 40s restarting failed container=pxc pod=cluster1-pxc-2_default(87cdf1a8-0fb3-4bc0-b50d-f66a0a73c087); ",
    "ready":2,
    "size":3,
    "status":"initializing"
  },
  "state":"initializing"
}
}

```

15.3.7 Pass custom my.cnf during the creation of Percona XtraDB Cluster

Description:

Create a custom config map containing the contents of the file my.cnf to be passed on to the Percona XtraDB Cluster containers when they are created

Kubectl Command:

```
$ kubectl create configmap cluster1-pxc3 --from-file=my.cnf
```

my.cnf (Contains mysql configuration):

```
[mysqld]
max_connections=250
```

URL:

```
https://$API_SERVER/api/v1/namespaces/default/configmaps
```

Authentication:

```
Authorization: Bearer $KUBE_TOKEN
```

cURL Request:

```
$ curl -k -v -XPOST "https://$API_SERVER/api/v1/namespaces/default/configmaps" \
-H "Accept: application/json" \
-H "Authorization: Bearer $KUBE_TOKEN" \
-d '{"apiVersion":"v1","data":{"my.cnf":"[mysqld]\nmax_connections=250\n"},"kind":"ConfigMap","metadata":{"creationTimestamp":null,"name":"cluster1-pxc3"}}' \
-H "Content-Type: application/json"
```

Request Body:

Example

```

{
  "apiVersion":"v1",
  "data":{
    "my.cnf":"[mysqld]\nmax_connections=250\n"
  },
  "kind":"ConfigMap",
  "metadata":{
    "creationTimestamp":null,
    "name":"cluster1-pxc3"
  }
}

```

Input:

1. data (Object {filename : contents(String, min-length:0)}): contains filenames to create in config map and its contents
2. metadata: name(String, min-length: 1) : contains name of the configmap
3. kind (String): type of object to create

Response:

Example

```
{
  "kind": "ConfigMap",
  "apiVersion": "v1",
  "metadata": {
    "name": "cluster1-pxc3",
    "namespace": "default",
    "selfLink": "/api/v1/namespaces/default/configmaps/cluster1-pxc3",
    "uid": "d92c7196-f399-4e20-abc7-b5de62c0691b",
    "resourceVersion": "85258",
    "creationTimestamp": "2020-05-28T14:19:41Z",
    "managedFields": [
      {
        "manager": "kubectl",
        "operation": "Update",
        "apiVersion": "v1",
        "time": "2020-05-28T14:19:41Z",
        "fieldsType": "FieldsV1",
        "fieldsV1": {
          "f:data": {
            ".": {
              },
              "f:my.cnf": {
                }
              }
            }
          }
        ]
      },
      {
        "data": {
          "my.cnf": ""
        }
      }
    ]
  }
}
```

15.3.8 Backup Percona XtraDB Cluster

Description:

Takes a backup of the Percona XtraDB Cluster containers data to be able to recover from disasters or make a roll-back later

Kubectl Command:

```
$ kubectl apply -f percona-xtradb-cluster-operator/deploy/backup/backup.yaml
```

URL:

```
https://$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusterbackups
```

Authentication:

```
Authorization: Bearer $KUBE_TOKEN
```

cURL Request:

```
$ curl -k -v -XPOST "https://$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusterbackups" \
  -H "Accept: application/json" \
  -H "Content-Type: application/json" \
  -d "@backup.json" -H "Authorization: Bearer $KUBE_TOKEN"
```

Request Body (backup.json):

Example

```
{
  "apiVersion": "pxc.percona.com/v1",
  "kind": "PerconaXtraDBClusterBackup",
  "metadata": {
    "name": "backup1"
  },
  "spec": {
    "pxcCluster": "cluster1",
    "storageName": "fs-pvc"
  }
}
```

Input:**1. metadata:**

name(String, min-length:1) : name of backup to create

1. spec:

1. pxcCluster(String, min-length:1) : `name of Percona XtraDB Cluster`
2. storageName(String, min-length:1) : `name of storage claim to use`

Response:**Example**

```
{
  "apiVersion": "pxc.percona.com/v1",
  "kind": "PerconaXtraDBClusterBackup",
  "metadata": {
    "creationTimestamp": "2020-05-27T23:56:33Z",
    "generation": 1,
    "managedFields": [
      {
        "apiVersion": "pxc.percona.com/v1",
        "fieldsType": "FieldsV1",
        "fieldsV1": {
          "f:spec": {
            ".": {
              "pxcCluster": {
                "f:pxcCluster": {
                  "f:storageName": {
                    "f:storageName": {}
                  }
                }
              }
            }
          }
        },
        "manager": "kubect1",
        "operation": "Update",
        "time": "2020-05-27T23:56:33Z"
      }
    ],
    "name": "backup1",
    "namespace": "default",
    "resourceVersion": "26024",
    "selfLink": "/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusterbackups/backup1",
    "uid": "95a354b1-e25b-40c3-8be4-388acca055fe"
  },
  "spec": {
    "pxcCluster": "cluster1",
    "storageName": "fs-pvc"
  }
}
```

15.3.9 Restore Percona XtraDB Cluster**Description:**

Restores Percona XtraDB Cluster data to an earlier version to recover from a problem or to make a roll-back

Kubect1 Command:

```
$ kubectl apply -f percona-xtradb-cluster-operator/deploy/backup/restore.yaml
```

URL:

```
https://$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusterrestores
```

Authentication:

```
Authorization: Bearer $KUBE_TOKEN
```

cURL Request:

```
$ curl -k -v -XPOST "https://$API_SERVER/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusterrestores" \
-H "Accept: application/json" \
-H "Content-Type: application/json" \
-d "@restore.json" \
-H "Authorization: Bearer $KUBE_TOKEN"
```

Request Body (restore.json):**Example**

```
{
  "apiVersion": "pxc.percona.com/v1",
  "kind": "PerconaXtraDBClusterRestore",
  "metadata": {
    "name": "restore1"
  },
  "spec": {
    "pxcCluster": "cluster1",
    "backupName": "backup1"
  }
}
```

Input:**1. metadata:**

name(String, min-length:1): name of restore to create

1. spec:

1. pxcCluster(String, min-length:1) : `name of Percona XtraDB Cluster`
2. backupName(String, min-length:1) : `name of backup to restore from`

Response:

Example

```
{
  "apiVersion": "pxc.percona.com/v1",
  "kind": "PerconaXtraDBClusterRestore",
  "metadata": {
    "creationTimestamp": "2020-05-27T23:59:41Z",
    "generation": 1,
    "managedFields": [
      {
        "apiVersion": "pxc.percona.com/v1",
        "fieldsType": "FieldsV1",
        "fieldsV1": {
          "f:spec": {
            ".": {
              },
              "f:backupName": {
                },
              "f:pxcCluster": {
                }
            }
          },
          "manager": "kubect1",
          "operation": "Update",
          "time": "2020-05-27T23:59:41Z"
        }
      ],
      "name": "restore1",
      "namespace": "default",
      "resourceVersion": "26682",
      "selfLink": "/apis/pxc.percona.com/v1/namespaces/default/perconaxtradbclusterrestores/restore1",
      "uid": "770c3471-be17-46fb-b0a6-e706685ab2fc"
    },
    "spec": {
      "backupName": "backup1",
      "pxcCluster": "cluster1"
    }
  }
}
```

Last update: 2022-12-04

15.4 Frequently Asked Questions

15.4.1 Why do we need to follow “the Kubernetes way” when Kubernetes was never intended to run databases?

As it is well known, the Kubernetes approach is targeted at stateless applications but provides ways to store state (in Persistent Volumes, etc.) if the application needs it. Generally, a stateless mode of operation is supposed to provide better safety, sustainability, and scalability, it makes the already-deployed components interchangeable. You can find more about substantial benefits brought by Kubernetes to databases in [this blog post](#).

The architecture of state-centric applications (like databases) should be composed in a right way to avoid crashes, data loss, or data inconsistencies during hardware failure. Percona Operator for MySQL provides out-of-the-box functionality to automate provisioning and management of highly available MySQL database clusters on Kubernetes.

15.4.2 How can I contact the developers?

The best place to discuss Percona Operator for MySQL based on Percona XtraDB Cluster with developers and other community members is the [community forum](#).

If you would like to report a bug, use the [Percona Operator for MySQL project in JIRA](#).

15.4.3 What is the difference between the Operator quickstart and advanced installation ways?

As you have noticed, the installation section of docs contains both quickstart and advanced installation guides.

The quickstart guide is simpler. It has fewer installation steps in favor of predefined default choices. Particularly, in advanced installation guides, you separately apply the Custom Resource Definition and Role-based Access Control configuration files with possible edits in them. At the same time, quickstart guides rely on the all-inclusive bundle configuration.

At another point, quickstart guides are related to specific platforms you are going to use (Minikube, Google Kubernetes Engine, etc.) and therefore include some additional steps needed for these platforms.

Generally, rely on the quickstart guide if you are a beginner user of the specific platform and/or you are new to the Percona Distribution for MySQL Operator as a whole.

15.4.4 Which versions of MySQL the Percona Operator for MySQL supports?

Percona Operator for MySQL based on Percona XtraDB Cluster provides a ready-to-use installation of the MySQL-based Percona XtraDB Cluster inside your Kubernetes installation. It works with both MySQL 8.0 and 5.7 branches, and the exact version is determined by the Docker image in use.

Percona-certified Docker images used by the Operator are listed [here](#). As you can see, both Percona XtraDB Cluster 8.0 and 5.7 are supported with the following recommended versions: 8.0.29-21.1 and 5.7.39-31.61. Three major numbers in the XtraDB Cluster version refer to the version of Percona Server in use. More details on the exact Percona Server version can be found in the release notes ([8.0](#), [5.7](#)).

15.4.5 How HAProxy is better than ProxySQL?

Percona Operator for MySQL based on Percona XtraDB Cluster supports both HAProxy and ProxySQL as a load balancer. HAProxy is turned on by default, but both solutions are similar in terms of their configuration and operation under the control of the Operator.

Still, they have technical differences. HAProxy is a general and widely used high availability, load balancing, and proxying solution for TCP and HTTP-based applications. ProxySQL provides similar functionality but is specific to MySQL clusters. As an SQL-aware solution, it is able to provide more tight internal integration with MySQL instances.

Both projects do a really good job with the Operator. The proxy choice should depend mostly on application-specific workload (including object-relational mapping), performance requirements, advanced routing and caching needs with one or another project, components already in use in the current infrastructure, and any other specific needs of the application.

15.4.6 How can I add create a directory on the node to use it as a local storage

You can [configure hostPath volume](#) to mount some existing file or directory from the node's filesystem into the Pod and use it as a local storage. The directory used for local storage should already exist in the node's filesystem. You can create it through the shell access to the node, with `mkdir` command, as all other directories. Alternatively you can create a Pod which will do this job. Let's suppose you are going to use `/var/run/data-dir` directory as your local storage, describing it in the `deploy/cr.yaml` configuration file as follows:

```
...
pxc:
  ...
  volumeSpec:
    hostPath:
      path: /var/run/data-dir
      type: Directory
  containerSecurityContext:
    privileged: false
  podSecurityContext:
    runAsUser: 1001
    runAsGroup: 1001
    supplementalGroups: [1001]
  nodeSelector:
    kubernetes.io/hostname: a.b.c
```

Create the yaml file (e.g. `mypod.yaml`), with the following contents:

```
apiVersion: v1
kind: Pod
metadata:
  name: hostpath-helper
spec:
  containers:
  - name: init
    image: busybox
    command: ["install", "-o", "1001", "-g", "1001", "-m", "755", "-d", "/mnt/data-dir"]
    volumeMounts:
    - name: root
      mountPath: /mnt
    securityContext:
      runAsUser: 0
  volumes:
  - name: root
    hostPath:
      path: /var/run
  restartPolicy: Never
  nodeSelector:
    kubernetes.io/hostname: a.b.c
```

Don't forget to apply it as usual:

```
$ kubectl apply -f mypod.yaml
```

15.4.7 How can I add custom sidecar containers to my cluster?

The Operator allows you to deploy additional (so-called *sidecar*) containers to the Pod. You can use this feature to run debugging tools, some specific monitoring solutions, etc. Add such sidecar container to the `deploy/cr.yaml` configuration file, specifying its name and image, and possibly a command to run:

```
spec:
  pxc:
    ...
    sidecars:
    - image: busybox
      command: ["/bin/sh"]
      args: ["-c", "while true; do echo echo $(date -u) 'test' >> /dev/null; sleep 5; done"]
      name: my-sidecar-1
    ...
```

You can add `sidecars` subsection to `pxc`, `haproxy`, and `proxysql` sections.

Note

Custom sidecar containers [can easily access other components of your cluster](#). Therefore they should be used carefully and by experienced users only.

Find more information on sidecar containers in the appropriate [documentation page](#).

15.4.8 How to get core dumps in case of the Percona XtraDB Cluster crash

In the Percona XtraDB Cluster crash case, gathering all possible information for enhanced diagnostics to be shared with Percona Support helps to solve an issue faster. One of such helpful artifacts is [core dump](#).

Percona XtraDB Cluster can create core dumps on crush [using libcoredumper](#). The Operator has this feature turned on by default. Core dumps are saved to DATADIR (`/var/lib/mysql/`). You can find appropriate core files in the following way (substitute `some-name-pxc-1` with the name of your Pod):

```
$ kubectl exec some-name-pxc-1 -c pxc -it -- sh -c 'ls -alh /var/lib/mysql/ | grep core'
-rw----- 1 mysql mysql 1.3G Jan 15 09:30 core.20210015093005
```

When identified, the appropriate core dump can be downloaded as follows:

```
$ kubectl cp some-name-pxc-1:/var/lib/mysql/core.20210015093005 /tmp/core.20210015093005
```

Note

It is useful to provide Build ID and Server Version in addition to core dump when Creating a support ticket. Both can be found from logs:

```
$ kubectl logs some-name-pxc-1 -c logs
```

```
[1] init-deploy-949.some-name-pxc-1.mysql-d-error.log: [1610702394.259356066, {"log"=>"09:19:54 UTC - mysqld got signal 11 ;"}]
[2] init-deploy-949.some-name-pxc-1.mysql-d-error.log: [1610702394.259356829, {"log"=>"Most likely, you have hit a bug, but this error can also be caused by malfunctioning hardware."}]
[3] init-deploy-949.some-name-pxc-1.mysql-d-error.log: [1610702394.259457282, {"log"=>"Build ID: 5a2199b1784b967a713a3bde8d996dc517c41adb"}]
[4] init-deploy-949.some-name-pxc-1.mysql-d-error.log: [1610702394.259465692, {"log"=>"Server Version: 8.0.21-12.1 Percona XtraDB Cluster (GPL), Release rel12, Revision 4d973e2, WSREP version 26.4.3, wsrep_26.4.3"}]
.....
```

15.4.9 How to choose between HAProxy and ProxySQL when configuring the cluster?

You can configure the Operator to use one of two different proxies, HAProxy (the default choice) and ProxySQL. Both solutions are fully supported by the Operator, but they have some differences in the architecture, which can make one of them more suitable than the other one in some use cases.

The main difference is that HAProxy operates in TCP mode as an [OSI level 4 proxy](#), while ProxySQL implements OSI level 7 proxy, and thus can provide some additional functionality like read/write split, firewalling and caching.

From the other side, utilizing HAProxy for the service is the easier way to go, and getting use of the ProxySQL level 7 specifics requires good understanding of Kubernetes and ProxySQL.

See more detailed functionality and performance comparison of using the Operator with both solutions in [this blog post](#).

15.4.10 Which additional access permissions are used by the Custom Resource validation webhook?

The `spec.enableCRValidationWebhook` key in the [deploy/cr.yaml](#) file enables or disables schema validation done by the Operator before applying `cr.yaml` file. This feature works only in [cluster-wide mode](#) due to access restrictions. It uses the following additional RBAC permissions:

```
- apiGroups:
  - admissionregistration.k8s.io
  resources:
  - validatingwebhookconfigurations
  verbs:
```

```
- get
- list
- watch
- create
- update
- patch
- delete
```

Last update: 2022-12-04

15.5 Old releases (documentation archive)

The documentation for the Operator version prior to 1.12.0 can be downloaded here in PDF format:

- [Percona Operator for MySQL based on Percona XtraDB Cluster 1.11.0](#)
- [Percona Distribution for MySQL Operator 1.10.0](#)
- [Percona Distribution for MySQL Operator 1.9.0](#)
- [Percona Kubernetes Operator for Percona XtraDB Cluster 1.8.0](#)
- [Percona Kubernetes Operator for Percona XtraDB Cluster 1.7.0](#)
- [Percona Kubernetes Operator for Percona XtraDB Cluster 1.6.0](#)
- [Percona Kubernetes Operator for Percona XtraDB Cluster 1.5.0](#)
- [Percona Kubernetes Operator for Percona XtraDB Cluster 1.4.0](#)
- [Percona Kubernetes Operator for Percona XtraDB Cluster 1.3.0](#)
- [Percona Kubernetes Operator for Percona XtraDB Cluster 1.2.0](#)
- [Percona Kubernetes Operator for Percona XtraDB Cluster 1.1.0](#)

Last update: 2022-12-07

16. Release Notes

16.1 Percona Operator for MySQL based on Percona XtraDB Cluster 1.12.0

- **Date**

December 7, 2022

- **Installation**

[Installing Percona Operator for MySQL based on Percona XtraDB Cluster](#)

16.1.1 Release Highlights

- [Azure Kubernetes Service \(AKS\)](#) is now officially supported platform, so developers and vendors of the solutions based on the Azure platform can take advantage of the official support from Percona or just use officially certified Percona Operator for MySQL images; also, [Azure Blob Storage can now be used for backups](#)
- This release also includes fixes to the following CVEs (Common Vulnerabilities and Exposures): [CVE-2021-20329](#) (potential injections in MongoDB Go Driver used HAProxy, which had no effect on Percona Operator for MySQL), and [CVE-2022-42898](#) (images used by the Operator suffering from the unauthenticated denial of service vulnerability). Users of previous Operator versions are advised to [upgrade](#) to version 1.12.0 which resolves this issue

16.1.2 New Features

- [K8SPXC-1043](#) and [K8SPXC-1005](#): Add support for the [Azure Kubernetes Service \(AKS\)](#) platform and allow [using Azure Blob Storage](#) for backups
- [K8SPXC-1010](#): Allow [using templates](#) to define `innodb_buffer_pool_size` auto-tuning based on container memory limits
- [K8SPXC-1082](#): New `ignoreAnnotations` and `ignoreLabels` Custom Resource options allow to list [specific annotations and labels](#) for Kubernetes Service objects, which the Operator should ignore (useful with various Kubernetes flavors which add annotations to the objects managed by the Operator)
- [K8SPXC-1120](#): Add [headless service](#) support for the restore Pod to [make it possible](#) restoring backups from a Persistent Volume on a tenant network (thanks to Zulh for contribution)
- [K8SPXC-1140](#): The Operator now [allows using SSL channel](#) for cross-site replication (thanks to Alvaro Aguilar-Tablada Espinosa for contribution)

16.1.3 Improvements

- [K8SPXC-1104](#): Starting from now, the Operator changed its API version to v1 instead of having a separate API version for each release. Three last API version are supported in addition to `v1`, which substantially reduces the size of Custom Resource Definition to prevent reaching the etcd limit
- [K8SPXC-955](#): Add Custom Resource options to set static IP-address for the [HAProxy](#) and [ProxySQL](#) LoadBalancers
- [K8SPXC-1032](#): Disable [automated upgrade](#) by default to prevent an unplanned downtime for user applications and to provide defaults more focused on strict user's control over the cluster
- [K8SPXC-1095](#): Process the SIGTERM signal to avoid unneeded lags in case of Percona XtraDB Cluster recovery or using the debug image to start up
- [K8SPXC-1113](#): Utilize dual password feature of MySQL 8 to avoid cluster restart when changing password of the `monitor` user
- [K8SPXC-1125](#): The Operator now does not attempt to start Percona Monitoring and Management (PMM) client sidecar if the corresponding secret does not contain the `pmmserver` or `pmmserverkey` key

- [K8SPXC-1153](#): Configuring the log structuring and leveling [is now supported](#) using the `LOG_STRUCTURED` and `LOG_LEVEL` environment variables. This reduces the information overload in logs, still leaving the possibility of getting more details when needed, for example, for debugging
- [K8SPXC-1123](#): Starting from now, installing the Operator for cluster-wide (multi-namespace) doesn't require to add Operator's own namespace to the list of watched namespaces (thanks to Bart Vercoulen for reporting this issue)
- [K8SPXC-1030](#): The new `delete-ssl` finalizer can now be used to automatically delete objects created for SSL (Secret, certificate, and issuer) in case of cluster deletion

16.1.4 Bugs Fixed

- [K8SPXC-1158](#): Fix [CVE-2022-42898](#) vulnerability found in MIT krb5, which made images used by the Operator vulnerable to DoS attacks
- [K8SPXC-1028](#): Fix a bug that prevented the Operator to automatically tune `innodb_buffer_pool_size` and `innodb_buffer_pool_chunk_size` variables
- [K8SPXC-1036](#): Fix the bug that caused Liveness Probe failure when XtraBackup was running and the `wsrep_sync_wait` option was set, making the instance to be rejected from the cluster
- [K8SPXC-1065](#): Fix a bug due to which, in a pair of scheduled backups close in time, the next backup could overwrite the previous one: bucket destination was made more unique by including seconds
- [K8SPXC-1059](#): Fix a bug due to which `pxc-monit` and `proxysql-monit` containers were printing passwords in their logs (thanks to zlcnju for contribution)
- [K8SPXC-1099](#): Fix CrashLoopBackOff error caused by incorrect (non-atomic) multi-user password change
- [K8SPXC-1100](#): Fix a bug that made it impossible to use slash characters in the monitor user's password
- [K8SPXC-1118](#): Fix a bug due to which the point-in-time recovery collector only reported warnings in logs when the gaps in binlogs were found. Starting from now, such backups are marked as not suitable for consistent PITR, and [restoring them with point-in-time recovery fails](#) without manual user's intervention
- [K8SPXC-1137](#): Fix a bug that prevented adding, deleting or updating ProxySQL Service labels/annotations except at the Service creation time
- [K8SPXC-1138](#): Fix a bug due to which not enough responsive scripts for readiness and liveness Probes could be the reason of killing the overloaded database Pods

16.1.5 Supported Platforms

The following platforms were tested and are officially supported by the Operator 1.12.0:

- [Google Kubernetes Engine \(GKE\)](#) 1.21 - 1.24
- [Amazon Elastic Container Service for Kubernetes \(EKS\)](#) 1.21 - 1.24
- [Azure Kubernetes Service \(AKS\)](#) 1.22 - 1.24
- [OpenShift](#) 4.10 - 4.11
- [Minikube](#) 1.28

This list only includes the platforms that the Percona Operators are specifically tested on as part of the release process. Other Kubernetes flavors and versions depend on the backward compatibility offered by Kubernetes itself.

Last update: 2022-12-07

16.2 Percona Operator for MySQL based on Percona XtraDB Cluster 1.11.0

- **Date**

June 3, 2022

- **Installation**

[Installing Percona Operator for MySQL based on Percona XtraDB Cluster](#)

16.2.1 Release Highlights

- With this release, the Operator turns to a simplified naming convention and changes its official name to **Percona Operator for MySQL based on Percona XtraDB Cluster**
- The new `backup.backoffLimit` Custom Resource option allows customizing the number of attempts the Operator should make for backup
- The OpenAPI schema is now generated for the Operator, which allows Kubernetes to validate Custom Resource and protects users from occasionally applying `deploy/cr.yaml` with syntax errors

16.2.2 New Features

- [K8SPXC-936](#): Allow modifying the init script via Custom Resource, which is useful for troubleshooting the Operator's issues
- [K8SPXC-758](#): Allow to [skip TLS verification for backup storage](#), useful for self-hosted S3-compatible storage with a self-signed certificate

16.2.3 Improvements

- [K8SPXC-947](#): Parametrize the number of attempt the Operator should make for backup backup through a [Custom Resource option](#)
- [K8SPXC-738](#): Allow to set service labels [for HAProxy](#) and [ProxySQL](#) in Custom Resource to enable various integrations with cloud providers or service meshes
- [K8SPXC-848](#): PMM container does not cause the crash of the whole database Pod if pmm-agent is not working properly
- [K8SPXC-625](#): Print the total number of binlogs and the number of remaining binlogs in the restore log while point-in-time recovery in progress
- [K8SPXC-920](#): Using the new [Percona XtraBackup Exponential Backoff feature](#) decreases the number of occasional unsuccessful backups due to more effective retries timing (Thanks to Dustin Falgout for reporting this issue)
- [K8SPXC-823](#): Make it possible [to use API Key](#) to authorize within Percona Monitoring and Management Server

16.2.4 Bugs Fixed

- [K8SPXC-985](#): Fix a bug that caused point-in-time recovery to fail due to incorrect binlog filtering logic
- [K8SPXC-899](#): Fix a bug due to which issued certificates didn't cover all hostnames, making `VERIFY_IDENTITY` client mode not working with HAProxy
- [K8SPXC-750](#): Fix a bug that prevented ProxySQL from connecting to Percona XtraDB Cluster after turning TLS off
- [K8SPXC-896](#): Fix a bug due to which the Operator was unable to create ssl-internal Secret if crash happened in the middle of a reconcile and restart (Thanks to srteam2020 for contribution)
- [K8SPXC-725](#) and [K8SPXC-763](#): Fix a bug due to which ProxySQL StatefulSet, and Services were mistakenly deleted by the Operator when reading stale ProxySQL or HAProxy information (Thanks to srteam2020 for contribution)
- [K8SPXC-957](#): Fix a bug due to which `pxc-db` Helm chart didn't support setting the `replicasServiceType` Custom Resource option (Thanks to Carlos Martell for reporting this issue)
- [K8SPXC-534](#): Fix a bug that caused some SQL queries to fail during the pxc StatefulSet update (Thanks to Sergiy Prykhodko for reporting this issue)

- [K8SPXC-1016](#): Fix a bug due to which an empty SSL secret name in Custom Resource caused the Operator to throw a misleading error message in the log
- [K8SPXC-994](#): Don't use root user in MySQL Pods to perform checks during cluster restoration, which may be helpful when restoring from non-Kubernetes environments
- [K8SPXC-961](#): Fix a bug due to which a user-defined sidecar container image in the Operator Pod could be treated as the `initImage` (Thanks to Carlos Martell for reporting this issue)
- [K8SPXC-934](#): Fix a bug due to which the cluster was not starting as Operator didn't create the users Secret if the `secretsName` option was absent in `cr.yaml`
- [K8SPXC-926](#): Fix a bug due to which failed Smart Update for one cluster in cluster-wide made the Operator unusable for other clusters
- [K8SPXC-900](#): Fix a bug where ProxySQL could not apply new configuration settings
- [K8SPXC-862](#): Fix a bug due to which changing resources as integer values without quotes in Custom Resource could lead to cluster getting stuck
- [K8SPXC-858](#): Fix a bug which could cause a single-node cluster to jump temporarily into the Error status during the upgrade
- [K8SPXC-814](#): Fix a bug when Custom Resource status was missing due to invalid variable setting in the manifest

16.2.5 Deprecation, Rename and Removal

- [K8SPXC-823](#): Password-based authorization to Percona Monitoring and Management Server is now deprecated and will be removed in future releases in favor of a token-based one. Password-based authorization was used by the Operator before this release to provide MySQL monitoring, but now using the API Key [is the recommended authorization method](#)

16.2.6 Supported Platforms

The following platforms were tested and are officially supported by the Operator 1.11.0:

- [OpenShift 4.7 - 4.10](#)
- [Google Kubernetes Engine \(GKE\) 1.20 - 1.23](#)
- [Amazon Elastic Container Service for Kubernetes \(EKS\) 1.20 - 1.22](#)
- [Minikube 1.23](#)

This list only includes the platforms that the Percona Operators are specifically tested on as part of the release process. Other Kubernetes flavors and versions depend on the backward compatibility offered by Kubernetes itself.

Last update: 2022-09-22

16.3 Percona Distribution for MySQL Operator 1.10.0

- **Date**

November 24, 2021

- **Installation**

For installation please refer to [the documentation page](#)

16.3.1 Release Highlights

- [Custom sidecar containers](#) allow users to customize Percona XtraDB Cluster and other Operator components without changing the container images. In this release, we enable even more customization, by allowing users to mount volumes into the sidecar containers.
- In this release, we put a lot of effort into fixing bugs that were reported by the community. We appreciate everyone who helped us with discovering these issues and contributed to the fixes.

16.3.2 New Features

- [K8SPXC-856](#): Mount volumes into sidecar containers to enable customization (Thanks to Sridhar L for contributing)

16.3.3 Improvements

- [K8SPXC-771](#): `spec.Backup.serviceAccount` and `spec.automountServiceAccountToken` Custom Resource options can now be used in the Helm chart (Thanks to Gerwin van de Steeg for reporting this issue)
- [K8SPXC-794](#): The `logrotate` command now doesn't use verbose mode to avoid flooding the log with rotate information
- [K8SPXC-793](#): Logs are now strictly following JSON specification to simplify parsing
- [K8SPXC-789](#): New `source_retry_count` and `source_connect_retry` options were added to tune source retries for replication between two clusters
- [K8SPXC-588](#): New `replicasServiceEnabled` option was added to allow disabling the Kubernetes Service for `haproxy-replicas`, which may be useful to avoid the unwanted forwarding of the application write requests to all Percona XtraDB Cluster instances
- [K8SPXC-822](#): Logrotate now doesn't rotate GRA logs (binlog events in ROW format representing the failed transaction) as ordinary log files, storing them for 7 days instead which gives additional time to debug the problem

16.3.4 Bugs Fixed

- [K8SPXC-761](#): Fixed a bug where HAProxy container was not setting explicit USER id, being incompatible with the `runAsNonRoot` security policy (Thanks to Henno Schooljan for reporting this issue)
- [K8SPXC-894](#): Fixed a bug where trailing white spaces in the `pmm-admin add` command caused reconcile loop on OpenShift
- [K8SPXC-831](#): Fixed a bug that made it possible to have a split-brain situation, when two nodes were starting their own cluster in case of a DNS failure
- [K8SPXC-796](#): Fixed a bug due to which S3 backup deletion didn't delete Pods attached to the backup job if the S3 finalizer was set (Thanks to Ben Langfeld for reporting this issue)
- [K8SPXC-876](#): Stopped using the `service.alpha.kubernetes.io/tolerate-unready-endpoints` deprecated Kubernetes option in the `$(cat $(clustername)-pxc-unready service annotation)` (Thanks to Antoine Habran for reporting this issue)
- [K8SPXC-842](#): Fixed a bug where backup finalizer didn't delete data from S3 if the backup path contained a folder inside of the S3 bucket (Thanks to 申祥瑞 for reporting this issue)
- [K8SPXC-812](#): Fix a bug due to which the Operator didn't support cert-manager versions since v0.14.0 (Thanks to Ben Langfeld for reporting this issue)

- [K8SPXC-762](#): Fix a bug due to which the validating webhook was not accepting scale operation in the Operator cluster-wide mode (Thanks to Henno Schooljan for reporting this issue)
- [K8SPXC-893](#): Fix a bug where HAProxy service failed during the config validation check if there was a resolution fail with one of the PXC addresses
- [K8SPXC-871](#): Fix a bug that prevented removing a Percona XtraDB Cluster manual backup for PVC storage
- [K8SPXC-851](#): Fixed a bug where changing replication user password didn't work
- [K8SPXC-850](#): Fixed a bug where the default weight value wasn't set for a host in a replication channel
- [K8SPXC-845](#): Fixed a bug where using malformed cr.yaml caused stuck cases in cluster deletion
- [K8SPXC-838](#): Fixed a bug due to which the Log Collector and PMM containers with unspecified memory and CPU requests were inheriting them from the PXC container
- [K8SPXC-824](#): Cluster may get into an unrecoverable state with incomplete full crash
- [K8SPXC-818](#): Fixed a bug which made Pods with a custom config inside a Secret or a ConfigMap not restarting at config update
- [K8SPXC-783](#): Fixed a bug where the root user was able to modify the monitor and clustercheck system users, making the possibility of cluster failure or misbehavior

16.3.5 Supported Platforms

The following platforms were tested and are officially supported by the Operator 1.10.0:

- OpenShift 4.7 - 4.9
- Google Kubernetes Engine (GKE) 1.19 - 1.22
- Amazon Elastic Kubernetes Service (EKS) 1.17 - 1.21
- Minikube 1.22

This list only includes the platforms that the Percona Operators are specifically tested on as part of the release process. Other Kubernetes flavors and versions depend on the backward compatibility offered by Kubernetes itself.

Last update: 2022-09-22

16.4 Percona Distribution for MySQL Operator 1.9.0

- **Date**

August 9, 2021

- **Installation**

For installation please refer to [the documentation page](#)

16.4.1 Release Highlights

- Starting from this release, the Operator changes its official name to **Percona Distribution for MySQL Operator**. This new name emphasizes gradual changes which incorporated a collection of Percona's solutions to run and operate Percona Server for MySQL and Percona XtraDB Cluster, available separately as [Percona Distribution for MySQL](#).
- Now you can [see HAProxy metrics](#) in your favorite Percona Monitoring and Management (PMM) dashboards automatically.
- The [cross-site replication](#) feature allows an asynchronous replication between two Percona XtraDB Clusters, including scenarios when one of the clusters is outside of the Kubernetes environment. The feature is intended for the following use cases:
 - provide migrations of your Percona XtraDB Cluster to Kubernetes or vice versa,
 - migrate regular MySQL database to Percona XtraDB Cluster under the Operator control, or carry on backward migration,
 - enable disaster recovery capability for your cluster deployment.

16.4.2 New Features

- [K8SPXC-657](#): Use Secrets to store custom configuration with sensitive data for [Percona XtraDB Cluster](#), [HAProxy](#), and [ProxySQL](#) Pods
- [K8SPXC-308](#): Implement Percona XtraDB Cluster [asynchronous replication](#) within the Operator
- [K8SPXC-688](#): Define [environment variables](#) in the Custom Resource to provide containers with additional customizations

16.4.3 Improvements

- [K8SPXC-673](#): HAProxy Pods now come with Percona Monitoring and Management integration and support
- [K8SPXC-791](#): Allow [stopping the restart-on-fail loop](#) for Percona XtraDB Cluster and Log Collector Pods without special debug images
- [K8SPXC-764](#): Unblock backups even if just a single instance is available by setting the `allowUnsafeConfigurations` flag to true
- [K8SPXC-765](#): Automatically delete custom configuration ConfigMaps if the variable in Custom Resource was unset (Thanks to Oleksandr Levchenkov for contributing)
- [K8SPXC-734](#): Simplify manual recovery by automatically getting Percona XtraDB Cluster namespace in the pxc container entrypoint script (Thanks to Michael Lin for contributing)
- [K8SPXC-656](#): `imagePullPolicy` is now set for init container as well to avoid pulling and simplifying deployments in air-gapped environments (Thanks to Herberto Graça for contributing)
- [K8SPXC-511](#): Secret object containing system users passwords is now deleted along with the Cluster if `delete-pxc-pvc` finalizer is enabled (Thanks to Matthias Baur for contributing)
- [K8SPXC-772](#): All Service objects now have Percona XtraDB Cluster labels attached to them to enable label selector usage
- [K8SPXC-731](#): It is now possible to see the overall progress of the provisioning of Percona XtraDB Cluster resources and dependent components in Custom Resource status
- [K8SPXC-730](#): Percona XtraDB Cluster resource statuses in Custom Resource output (e.g. returned by `kubectl get pxc` command) have been improved and now provide more precise reporting
- [K8SPXC-697](#): Add namespace support in the `copy-backup` script

- [K8SPXC-321](#), [K8SPXC-556](#), [K8SPXC-568](#): Restrict the minimal number of ProxySQL and HAProxy Pods and the maximal number of Percona XtraDB Cluster Pods if the unsafe flag is not set
- [K8SPXC-554](#): Reduced the number of various etcd and k8s object updates from the Operator to minimize the pressure on the Kubernetes cluster
- [K8SPXC-421](#): It is now possible to [use X Plugin](#) with Percona XtraDB Cluster Pods

16.4.4 Known Issues and Limitations

- [K8SPXC-835](#): ProxySQL will fail to start on a Replica Percona XtraDB Cluster for cross-site replication in this release

16.4.5 Bugs Fixed

- [K8SPXC-757](#): Fixed a bug where manual crash recovery interfered with auto recovery functionality even with the `auto_recovery` flag set to false
- [K8SPXC-706](#): TLS certificates [renewal by a cert-manager was failing](#) (Thanks to Jeff Andrews for reporting this issue)
- [K8SPXC-785](#): Fixed a bug where backup to S3 was producing false-positive error messages even if backup was successful
- [K8SPXC-642](#): Fixed a bug where PodDisruptionBudget was blocking the upgrade of HAProxy (Thanks to Davi S Evangelista for reporting this issue)
- [K8SPXC-585](#): Fixed a bug where the Operator got stuck if the wrong user credentials were set in the Secret object (Thanks to Sergiy Prykhodko for reporting this issue)
- [K8SPXC-756](#): Fixed a bug where the Operator was scheduling backups even when the cluster was paused (Thanks to Dmytro for reporting this issue)
- [K8SPXC-813](#): Fixed a bug where backup restore didn't return error on incorrect AWS credentials
- [K8SPXC-805](#): Fixed a bug that made pxc-backups object deletion hang if the Operator couldn't list objects from the S3 bucket (e.g. due to wrong S3 credentials)
- [K8SPXC-787](#): Fixed the "initializing" status of ready clusters caused by the xtrabackup user password change
- [K8SPXC-775](#): Fixed a bug where errors in custom mysqld config settings were not detected by the Operator if the config was modified after the initial cluster was created
- [K8SPXC-767](#): Fixed a bug where on-demand backup hung up if created while the cluster was in the "initializing" state
- [K8SPXC-726](#): Fixed a bug where the `delete-s3-backup` finalizer prevented deleting a backup stored on Persistent Volume
- [K8SPXC-682](#): Fixed auto-tuning feature setting wrong `innodb_buffer_pool_size` value in some cases

Last update: 2022-09-22

16.5 Percona Kubernetes Operator for Percona XtraDB Cluster 1.8.0

- **Date**

April 26, 2021

- **Installation**

[Installing Percona Kubernetes Operator for Percona XtraDB Cluster](#)

16.5.1 Release Highlights

- It is now [possible](#) to use `kubectl scale` command to scale Percona XtraDB Cluster horizontally (add or remove Replica Set instances). You can also use [Horizontal Pod Autoscaler](#) which will scale your database cluster based on various metrics, such as CPU utilization.
- Support for [custom sidecar containers](#). The Operator makes it possible now to deploy additional (sidecar) containers to the Pod. This feature can be useful to run debugging tools or some specific monitoring solutions, etc. Sidecar containers can be added to [pxc](#), [haproxy](#), and [proxysql](#) sections of the `deploy/cr.yaml` configuration file.

16.5.2 New Features

- [K8SPXC-528](#): Support for [custom sidecar containers](#) to extend the Operator capabilities
- [K8SPXC-647](#): Allow the cluster [scale in and scale out](#) with the `kubectl scale` command or Horizontal Pod Autoscaler
- [K8SPXC-643](#): Operator can now automatically recover Percona XtraDB Cluster after the [network partitioning](#)

16.5.3 Improvements

- [K8SPXC-442](#): The Operator can now automatically remove old backups from S3 storage if the retention period is set (thanks to Davi S Evangelista for reporting this issue)
- [K8SPXC-697](#): Add namespace support in the [script used to copy backups](#) from remote storage to a local machine
- [K8SPXC-627](#): Point-in-time recovery uploader now chooses the Pod with the oldest binary log in the cluster to ensure log consistency
- [K8SPXC-618](#): Add debug symbols from the [percona-xtradb-cluster-server-debuginfo](#) package to the Percona XtraDB Cluster debug docker image to simplify troubleshooting
- [K8SPXC-599](#): It is now possible to [recover](#) databases up to a specific transaction with the Point-in-time Recovery feature. Previously the user could only recover to specific date and time
- [K8SPXC-598](#): Point-in-time recovery feature now works with compressed backups
- [K8SPXC-536](#): It is now possible to explicitly set the version of Percona XtraDB Cluster for newly provisioned clusters. Before that, all new clusters were started with the latest PXC version if Version Service was enabled
- [K8SPXC-522](#): Add support for the `runtimeClassName` Kubernetes feature for selecting the container runtime
- [K8SPXC-519](#), [K8SPXC-558](#), and [K8SPXC-637](#): Various improvements of Operator log messages

16.5.4 Known Issues and Limitations

- [K8SPXC-701](#): Scheduled backups are not compatible with Kubernetes 1.20 in cluster-wide mode.

16.5.5 Bugs Fixed

- [K8SPXC-654](#): Use MySQL administrative port for Kubernetes liveness/readiness probes to avoid false positive failures
- [K8SPXC-614](#), [K8SPXC-619](#), [K8SPXC-545](#), [K8SPXC-641](#), [K8SPXC-576](#): Fix multiple bugs due to which changes of various objects in `deploy/cr.yaml` were not applied to the running cluster (thanks to Sergiy Prykhodko for reporting some of these issues)

- [K8SPXC-596](#): Fix a bug due to which liveness probe for `pxc` container could cause zombie processes
 - [K8SPXC-632](#): Fix a bug preventing point-in-time recovery when multiple clusters were uploading binary logs to a single S3 bucket
 - [K8SPXC-573](#): Fix a bug that prevented using special characters in XtraBackup password (thanks to Gertjan Bijl for reporting this issue)
 - [K8SPXC-571](#): Fix a bug where Percona XtraDB Cluster went into a desynced state at backup job crash (Thanks to Dimitrij Hilt for reporting this issue)
 - [K8SPXC-430](#): Galera Arbitrator used for backups does not break the cluster anymore in various cases
 - [K8SPXC-684](#): Fix a bug due to which point-in-time recovery backup didn't allow specifying the `endpointUrl` for Amazon S3 storage
 - [K8SPXC-681](#): Fix operator crash which occurred when non-existing storage name was specified for point-in-time recovery
 - [K8SPXC-638](#): Fix unneeded delay in showing logs with the `kubectl logs` command for the logs container
 - [K8SPXC-609](#): Fix frequent HAProxy service NodePort updates which were causing issues with load balancers
 - [K8SPXC-542](#): Fix a bug due to which backups were taken only for one cluster out of many controlled by one Operator
 - [CLOUD-611](#): Stop using the already deprecated runtime/scheme package (Thanks to Jerome Küttner for reporting this issue)
-

Last update: 2022-09-22

16.6 Percona Kubernetes Operator for Percona XtraDB Cluster 1.7.0

- **Date**

February 2, 2021

- **Installation**

[Installing Percona Kubernetes Operator for Percona XtraDB Cluster](#)

16.6.1 New Features

- **K8SPXC-530**: Add support for [point-in-time recovery](#)
- **K8SPXC-564**: PXC cluster will now recover automatically from a full crash when Pods are stuck in CrashLoopBackOff status
- **K8SPXC-497**: Official support for [Percona Monitoring and Management \(PMM\) v.2](#)

NOTE: Monitoring with PMM v.1 configured according to the [unofficial instruction](#) will not work after the upgrade. Please switch to PMM v.2.

16.6.2 Improvements

- **K8SPXC-485**: [Percona XtraDB Cluster Pod logs are now stored on Persistent Volumes](#). Users can debug the issues even after the Pod restart
- **K8SPXC-389**: User can now change ServiceType for HAProxy replicas Kubernetes service
- **K8SPXC-546**: Reduce the number of ConfigMap object updates from the Operator to improve performance of the Kubernetes cluster
- **K8SPXC-553**: Change default configuration of ProxySQL to WRITERS_ARE_READERS=yes so Percona XtraDB Cluster continues operating with a single node left
- **K8SPXC-512**: User can now limit cluster-wide Operator access to specific namespaces (Thanks to user mgar for contribution)
- **K8SPXC-490**: Improve error message when not enough memory is set for auto-tuning
- **K8SPXC-312**: Add schema validation for Custom Resource. Now `cr.yaml` is validated by a WebHook for syntax typos before being applied. It works only in cluster-wide mode due to access restrictions
- **K8SPXC-510**: Percona XtraDB Cluster operator can now be [deployed through RedHat Marketplace](#)
- **K8SPXC-543**: Check HAProxy custom configuration for syntax errors before applying it to avoid Pod getting stuck in CrashLoopBackOff status (Thanks to user pservit for reporting this issue)

16.6.3 Bugs Fixed

- **K8SPXC-544**: Add a liveness probe for HAProxy so it is not stuck and automatically restarted when crashed (Thanks to user pservit for reporting this issue)
- **K8SPXC-500**: Fix a bug that prevented creating a backup in cluster-wide mode if default `cr.yaml` is used (Thanks to user michael.lin1 for reporting this issue)
- **K8SPXC-491**: Fix a bug due to which compressed backups didn't work with the Operator (Thanks to user dejw for reporting this issue)
- **K8SPXC-570**: Fix a bug causing backups to fail with some S3-compatible storages (Thanks to user dimitrij for reporting this issue)
- **K8SPXC-517**: Fix a bug causing Operator crash if Custom Resource backup section is missing (Thanks to user daemonmv for reporting this issue)
- **K8SPXC-253**: Fix a bug preventing rolling out Custom Resource changes (Thanks to user bitsbeats for reporting this issue)
- **K8SPXC-552**: Fix a bug when HAProxy secrets cannot be updated by the user

- [K8SPXC-551](#): Fix a bug due to which cluster was not initialized when the password had an end of line symbol in `secret.yaml`
 - [K8SPXC-526](#): Fix a bug due to which not all clusters managed by the Operator were upgraded by the automatic update
 - [K8SPXC-523](#): Fix a bug putting cluster into unhealthy status after the clustercheck secret changed
 - [K8SPXC-521](#): Fix automatic upgrade job repeatedly looking for an already removed cluster
 - [K8SPXC-520](#): Fix Smart update in cluster-wide mode adding version service check job repeatedly instead of doing it only once
 - [K8SPXC-463](#): Fix a bug due to which `wsrep_recovery` log was unavailable after the Pod restart
 - [K8SPXC-424](#): Fix a bug due to which HAProxy health-check spammed in logs, making them hardly unreadable
 - [K8SPXC-379](#): Fix a bug due to which the Operator user credentials were not added into internal secrets when upgrading from 1.4.0 (Thanks to user `pservit` for reporting this issue)
-

Last update: 2022-09-22

16.7 Percona Kubernetes Operator for Percona XtraDB Cluster 1.6.0

- **Date**

October 9, 2020

- **Installation**

[Installing Percona Kubernetes Operator for Percona XtraDB Cluster](#)

16.7.1 New Features

- [K8SPXC-394](#): Support of “cluster-wide” mode for Percona XtraDB Cluster Operator
- [K8SPXC-416](#): Support of the proxy-protocol in HAProxy (to use this feature, you should have a Percona XtraDB Cluster image version 8.0.21 or newer)
- [K8SPXC-429](#): A possibility to restore backups to a new Kubernetes-based environment with no existing Percona XtraDB Cluster Custom Resource
- [K8SPXC-343](#): Helm chart officially provided with the Operator

16.7.2 Improvements

- [K8SPXC-144](#): Allow adding ProxySQL configuration options
- [K8SPXC-398](#): New `crVersion` key in `deploy/cr.yaml` to indicate the API version that the Custom Resource corresponds to (thanks to user [mike.saah](#) for contribution)
- [K8SPXC-474](#): The init container now has the same resource requests as the main container of a correspondent Pod (thanks to user [yann.leenhardt](#) for contribution)
- [K8SPXC-372](#): Support new versions of cert-manager by the Operator (thanks to user [rf_enigm](#) for contribution)
- [K8SPXC-317](#): Possibility to configure the `imagePullPolicy` Operator option (thanks to user [imranrazakhan](#) for contribution)
- [K8SPXC-462](#): Add readiness probe for HAProxy
- [K8SPXC-411](#): Extend cert-manager configuration to add additional domains (multiple SAN) to a certificate
- [K8SPXC-375](#): Improve HAProxy behavior in case of switching writer node to a new one and back
- [K8SPXC-368](#): Autoupdate system users by changing the appropriate Secret name

16.7.3 Known Issues and Limitations

- OpenShift 3.11 requires additional configuration for the correct HAProxy operation: the feature gate `PodShareProcessNamespace` should be set to `true`. If getting it enabled is not possible, we recommend using ProxySQL instead of HAProxy with OpenShift 3.11. Other OpenShift and Kubernetes versions are not affected.
- [K8SPXC-491](#): Compressed backups are not compatible with the Operator 1.6.0 (`percona/percona-xtradb-cluster-operator:1.5.0-pxc8.0-backup` or `percona/percona-xtradb-cluster-operator:1.5.0-pxc5.7-backup` image can be used as a workaround if needed).

16.7.4 Bugs Fixed

- [K8SPXC-431](#): HAProxy unable to start on OpenShift with the default `cr.yaml` file
- [K8SPXC-408](#): Insufficient `MAX_USER_CONNECTIONS=10` for ProxySQL monitor user (increased to 100)
- [K8SPXC-391](#): HAProxy and PMM cannot be enabled at the same time (thanks to user [rf_enigm](#) for reporting this issue)
- [K8SPXC-406](#): Second node (XXX-pxc-1) always selected as a donor (thanks to user [pservit](#) for reporting this issue)
- [K8SPXC-390](#): Crash on missing HAProxy PodDisruptionBudget
- [K8SPXC-355](#): Counterintuitive YYYY-DD-MM dates in the S3 backup folder names (thanks to user [graham-web](#) for contribution)

- [K8SPXC-305](#): ProxySQL not working in case of passwords with a % symbol in the Secrets object (thanks to user ben.wilson for reporting this issue)
 - [K8SPXC-278](#): ProxySQL never getting ready status in some environments after the cluster launch due to the `proxysql-monitor` Pod crash (thanks to user lots0logs for contribution)
 - [K8SPXC-274](#): The 1.2.0 -> 1.3.0 -> 1.4.0 upgrade path not working (thanks to user martin.atroo for reporting this issue)
 - [K8SPXC-476](#): SmartUpdate failing to fetch version from Version Service in case of incorrectly formatted Percona XtraDB Cluster patch version higher than the last known one
 - [K8SPXC-454](#): After the cluster creation, pxc-0 Pod restarting due to Operator not waiting for cert-manager to issue requested certificates (thanks to user mike.saah for reporting this issue)
 - [K8SPXC-450](#): TLS annotations causing unnecessary HAProxy Pod restarts
 - [K8SPXC-443](#) and [K8SPXC-456](#): The outdated version service endpoint URL (fix with preserving backward compatibility)
 - [K8SPXC-435](#): MySQL root password visible through `kubectl logs`
 - [K8SPXC-426](#): mysqld recovery logs not logged to file and not available through `kubectl logs`
 - [K8SPXC-423](#): HAProxy not refreshing IP addresses even when the node gets a different address
 - [K8SPXC-419](#): Percona XtraDB Cluster incremental state transfers not taken into account by readiness/liveness checks
 - [K8SPXC-418](#): HAProxy not routing traffic for 1 donor, 2 joiners
 - [K8SPXC-417](#): Cert-manager not compatible with Kubernetes versions below v1.15 due to unnecessarily high API version demand
 - [K8SPXC-384](#): Debug images were not fully functional for the latest version of the Operator because of having no infinity loop
 - [K8SPXC-383](#): DNS warnings in PXC Pods when using HAProxy
 - [K8SPXC-364](#): Smart Updates showing empty "from" versions for non-PXC objects in logs
 - [K8SPXC-379](#): The Operator user credentials not added into internal secrets when upgrading from 1.4.0 (thanks to user pservit for reporting this issue)
-

Last update: 2022-09-22

16.8 Percona Kubernetes Operator for Percona XtraDB Cluster 1.5.0

- **Date**

July 21, 2020

- **Installation**

[Installing Percona Kubernetes Operator for Percona XtraDB Cluster](#)

16.8.1 New Features

- [K8SPXC-298](#): Automatic synchronization of MySQL users with ProxySQL
- [K8SPXC-294](#): HAProxy Support
- [K8SPXC-284](#): Fully automated minor version updates (Smart Update)
- [K8SPXC-257](#): Update Reader members before Writer member at cluster upgrades
- [K8SPXC-256](#): Support multiple PXC minor versions by the Operator

16.8.2 Improvements

- [K8SPXC-290](#): Extend usable backup schedule syntax to include lists of values
- [K8SPXC-309](#): Quickstart Guide on Google Kubernetes Engine (GKE) - [link](#)
- [K8SPXC-288](#): Quickstart Guide on Amazon Elastic Kubernetes Service (EKS) - [link](#)
- [K8SPXC-280](#): Support XtraBackup compression
- [K8SPXC-279](#): Use SYSTEM_USER privilege for system users on PXC 8.0
- [K8SPXC-277](#): Install GDB in PXC images
- [K8SPXC-276](#): Pod-0 should be selected as Writer if possible
- [K8SPXC-252](#): Automatically manage system users for MySQL and ProxySQL on password rotation via Secret
- [K8SPXC-242](#): Improve internal backup implementation for better stability with PXC 8.0
- [CLOUD-404](#): Support of loadBalancerSourceRanges for LoadBalancer Services
- [CLOUD-556](#): Kubernetes 1.17 added to the list of supported platforms

16.8.3 Bugs Fixed

- [K8SPXC-327](#): CrashloopBackOff if PXC 8.0 Pod restarts in the middle of SST
- [K8SPXC-291](#): PXC Restore failure with “The node was low on resource: ephemeral-storage” error (Thanks to user rjeka for reporting this issue)
- [K8SPXC-270](#): Restore job wiping data from the original backup’s cluster when restoring to another cluster in the same namespace
- [K8SPXC-352](#): Backup cronjob not scheduled in some Kubernetes environments (Thanks to user msavchenko for reporting this issue)
- [K8SPXC-275](#): Outdated documentation on the Operator updates (Thanks to user martin.atroo for reporting this issue)
- [K8SPXC-347](#): XtraBackup failure after uploading a backup, causing the backup process restart in some cases (Thanks to user connde for reporting this issue)
- [K8SPXC-373](#): Pod not cleaning up the SST tmp dir on start
- [K8SPXC-326](#): Changes in TLS Secrets not triggering PXC restart if AllowUnsafeConfig enabled
- [K8SPXC-323](#): Missing `tar` utility in the PXC node docker image
- [CLOUD-531](#): Wrong usage of `strings.TrimSpace` when processing `apiVersion`
- [CLOUD-474](#): Cluster creation not failing if wrong resources are set

Last update: 2022-09-22

16.9 Percona Kubernetes Operator for Percona XtraDB Cluster 1.4.0

- **Date**

April 29, 2020

- **Installation**

[Installing Percona Kubernetes Operator for Percona XtraDB Cluster](#)

16.9.1 New Features

- [K8SPXC-172](#): Full data-at-rest encryption available in PXC 8.0 is now supported by the Operator. This feature is implemented with the help of the `keyring_vault` plugin which ships with PXC 8.0. By utilizing [Vault](#) we enable our customers to follow best practices with encryption in their environment.
- [K8SPXC-125](#): Percona XtraDB Cluster 8.0 is now supported
- [K8SPXC-95](#): Amazon Elastic Container Service for Kubernetes (EKS) was added to the list of the officially supported platforms
- The OpenShift Container Platform 4.3 is now supported

16.9.2 Improvements

- [K8SPXC-262](#): The Operator allows setting ephemeral-storage requests and limits on all Pods
- [K8SPXC-221](#): The Operator now updates observedGeneration status message to allow better monitoring of the cluster rollout or backup/restore process
- [K8SPXC-213](#): A special [PXC debug image](#) is now available. It avoids restarting on fail and contains additional tools useful for debugging
- [K8SPXC-100](#): The Operator now implements the crash tolerance on the one member crash. The implementation is based on starting Pods with `mysqld --wsrep_recover` command if there was no graceful shutdown

16.9.3 Bugs Fixed

- [K8SPXC-153](#): S3 protocol credentials were not masked in logs during the PXC backup & restore process
- [K8SPXC-222](#): The Operator got caught in reconciliation error in case of the erroneous/absent API version in the `deploy/cr.yaml` file
- [K8SPXC-261](#): ProxySQL logs were showing the root password
- [K8SPXC-220](#): The inability to update or delete existing CRD was possible because of too large records in etcd, resulting in “request is too large” errors. Only 20 last status changes are now stored in etcd to avoid this problem.
- [K8SPXC-52](#): The Operator produced an unclear error message in case of fail caused by the absent or malformed pxc section in the `deploy/cr.yaml` file
- [K8SPXC-269](#): The `copy-backup.sh` script didn't work correctly in case of an existing secret with the `AWS_ACCESS_KEY_ID/AWS_SECRET_ACCESS_KEY` credentials and prevented users from copying backups (e.g. to a local machine)
- [K8SPXC-263](#): The `kubectl get pxc` command was unable to show the correct ProxySQL external endpoint
- [K8SPXC-219](#): PXC Helm charts were incompatible with the version 3 of the Helm package manager
- [K8SPXC-40](#): The cluster was unable to reach “ready” status in case if `ProxySQL.Enabled` field was set to `false`
- [K8SPXC-34](#): Change of the `proxysql.servicetype` field was not detected by the Operator and thus had no effect

Last update: 2022-09-22

16.10 Percona Kubernetes Operator for Percona XtraDB Cluster 1.3.0

Percona announces the *Percona Kubernetes Operator for Percona XtraDB Cluster* 1.3.0 release on January 6, 2020. This release is now the current GA release in the 1.3 series. [Install the Kubernetes Operator for Percona XtraDB Cluster by following the instructions.](#)

The Percona Kubernetes Operator for Percona XtraDB Cluster automates the lifecycle and provides a consistent Percona XtraDB Cluster instance. The Operator can be used to create a Percona XtraDB Cluster, or scale an existing Cluster and contains the necessary Kubernetes settings.

The Operator simplifies the deployment and management of the [Percona XtraDB Cluster](#) in Kubernetes-based environments. It extends the Kubernetes API with a new custom resource for deploying, configuring and managing the application through the whole life cycle.

The Operator source code is available [in our Github repository](#). All of Percona's software is open-source and free.

New features and improvements:

- **CLOUD-412:** Auto-Tuning of the MySQL Parameters based on Pod memory resources was implemented in the case of Percona XtraDB Cluster Pod limits (or at least Pod requests) specified in the cr.yaml file.
- **CLOUD-411:** Now the user can adjust securityContext, replacing the automatically generated securityContext with the customized one.
- **CLOUD-394:** The Percona XtraDB Cluster, ProxySQL, and backup images size decrease by 40-60% was achieved by removing unnecessary dependencies and modules to reduce the cluster deployment time.
- **CLOUD-390:** Helm chart for Percona Monitoring and Management (PMM) 2.0 has been provided.
- **CLOUD-383:** Affinity constraints and tolerations were added to the backup Pod
- **CLOUD-430:** Image URL in the CronJob Pod template is automatically updated when the Operator detects changed backup image URL

Fixed bugs:

- **CLOUD-462:** Resource requests/limits were set not for all containers in a ProxySQL Pod
- **CLOUD-437:** Percona Monitoring and Management Client was taking resources definition from the Percona XtraDB Cluster despite having much lower need in resources, particularly lower memory footprint.
- **CLOUD-434:** Restoring Percona XtraDB Cluster was failing on the OpenShift platform with customized security settings
- **CLOUD-399:** The iputils package was added to the backup docker image to provide backup jobs with the ping command for a better network connection handling
- **CLOUD-393:** The Operator generated various StatefulSets in the first reconciliation cycle and in all subsequent reconciliation cycles, causing Kubernetes to trigger an unnecessary ProxySQL restart once during the cluster creation.
- **CLOUD-376:** A long-running SST caused the liveness probe check to fail it's grace period timeout, resulting in an unrecoverable failure
- **CLOUD-243:** Using MYSQL_ROOT_PASSWORD with special characters in a ProxySQL docker image was breaking the entrypoint initialization process

[Percona XtraDB Cluster](#) is an open source, cost-effective and robust clustering solution for businesses. It integrates Percona Server for MySQL with the Galera replication library to produce a highly-available and scalable MySQL® cluster complete with synchronous multi-primary replication, zero data loss and automatic node provisioning using Percona XtraBackup.

Help us improve our software quality by reporting any bugs you encounter using [our bug tracking system](#).

Last update: 2022-09-22

16.11 Percona Kubernetes Operator for Percona XtraDB Cluster 1.2.0

Percona announces the *Percona Kubernetes Operator for Percona XtraDB Cluster 1.2.0* release on September 20, 2019. This release is now the current GA release in the 1.2 series. [Install the Kubernetes Operator for Percona XtraDB Cluster by following the instructions.](#)

The Percona Kubernetes Operator for Percona XtraDB Cluster automates the lifecycle and provides a consistent Percona XtraDB Cluster instance. The Operator can be used to create a Percona XtraDB Cluster, or scale an existing Cluster and contains the necessary Kubernetes settings.

The Operator simplifies the deployment and management of the [Percona XtraDB Cluster](#) in Kubernetes-based environments. It extends the Kubernetes API with a new custom resource for deploying, configuring and managing the application through the whole life cycle.

The Operator source code is available [in our Github repository](#). All of Percona's software is open-source and free.

New features and improvements:

- A [Service Broker](#) was implemented for the Operator, allowing a user to deploy Percona XtraDB Cluster on the OpenShift Platform, configuring it with a standard GUI, following the Open Service Broker API.
- Now the Operator supports [Percona Monitoring and Management 2](#), which means being able to detect and register to PMM Server of both 1.x and 2.0 versions.
- A `NodeSelector` constraint is now supported for the backups, which allows using backup storage accessible to a limited set of nodes only (contributed by [Chen Min](#)).
- The resource constraint values were refined for all containers to eliminate the possibility of an out of memory error.
- Now it is possible to set the `schedulerName` option in the operator parameters. This allows using storage which depends on a custom scheduler, or a cloud provider which optimizes scheduling to run workloads in a cost-effective way (contributed by [Smaïne Kahlouch](#)).
- A bug was fixed, which made cluster status oscillate between "initializing" and "ready" after an update.
- A 90 second startup delay which took place on freshly deployed Percona XtraDB Cluster was eliminated.

[Percona XtraDB Cluster](#) is an open source, cost-effective and robust clustering solution for businesses. It integrates Percona Server for MySQL with the Galera replication library to produce a highly-available and scalable MySQL® cluster complete with synchronous multi-primary replication, zero data loss and automatic node provisioning using Percona XtraBackup.

Help us improve our software quality by reporting any bugs you encounter using [our bug tracking system](#).

Last update: 2022-09-22

16.12 Percona Kubernetes Operator for Percona XtraDB Cluster 1.1.0

Percona announces the general availability of *Percona Kubernetes Operator for Percona XtraDB Cluster 1.1.0* on July 15, 2019. This release is now the current GA release in the 1.1 series. [Install the Kubernetes Operator for Percona XtraDB Cluster by following the instructions.](#)

The Percona Kubernetes Operator for Percona XtraDB Cluster automates the lifecycle and provides a consistent Percona XtraDB Cluster instance. The Operator can be used to create a Percona XtraDB Cluster, or scale an existing Cluster and contains the necessary Kubernetes settings.

The Operator simplifies the deployment and management of the [Percona XtraDB Cluster](#) in Kubernetes-based environments. It extends the Kubernetes API with a new custom resource for deploying, configuring and managing the application through the whole life cycle.

The Operator source code is available [in our Github repository](#). All of Percona's software is open-source and free.

New features and improvements:

- Now the Percona Kubernetes Operator [allows upgrading](#) Percona XtraDB Cluster to newer versions, either in semi-automatic or in manual mode.
- Also, two modes are implemented for updating the Percona XtraDB Cluster `my.cnf` configuration file: in *automatic configuration update* mode Percona XtraDB Cluster Pods are immediately re-created to populate changed options from the Operator YAML file, while in *manual mode* changes are held until Percona XtraDB Cluster Pods are re-created manually.
- A separate service account is now used by the Operator's containers which need special privileges, and all other Pods run on default service account with limited permissions.
- [User secrets](#) are now generated automatically if don't exist: this feature especially helps reduce work in repeated development environment testing and reduces the chance of accidentally pushing predefined development passwords to production environments.
- The Operator [is now able to generate TLS certificates itself](#) which removes the need in manual certificate generation.
- The list of officially supported platforms now includes [Minikube](#), which provides an easy way to test the Operator locally on your own machine before deploying it on a cloud.
- Also, Google Kubernetes Engine 1.14 and OpenShift Platform 4.1 are now supported.

[Percona XtraDB Cluster](#) is an open source, cost-effective and robust clustering solution for businesses. It integrates Percona Server for MySQL with the Galera replication library to produce a highly-available and scalable MySQL® cluster complete with synchronous multi-primary replication, zero data loss and automatic node provisioning using Percona XtraBackup.

Help us improve our software quality by reporting any bugs you encounter using [our bug tracking system](#).

Last update: 2022-09-22

16.13 Percona Kubernetes Operator for Percona XtraDB Cluster 1.0.0

Percona announces the general availability of *Percona Kubernetes Operator for Percona XtraDB Cluster 1.0.0* on May 29, 2019. This release is now the current GA release in the 1.0 series. [Install the Kubernetes Operator for Percona XtraDB Cluster by following the instructions](#). Please see the [GA release announcement](#). All of Percona's software is open-source and free.

The Percona Kubernetes Operator for Percona XtraDB Cluster automates the lifecycle and provides a consistent Percona XtraDB Cluster instance. The Operator can be used to create a Percona XtraDB Cluster, or scale an existing Cluster and contains the necessary Kubernetes settings.

The Percona Kubernetes Operators are based on best practices for configuration and setup of the Percona XtraDB Cluster. The Operator provides a consistent way to package, deploy, manage, and perform a backup and a restore for a Kubernetes application. Operators deliver automation advantages in cloud-native applications.

The advantages are the following:

- * Deploy a Percona XtraDB Cluster environment with no single point of failure and environment can span multiple availability zones (AZs).
- * Deployment takes about six minutes with the default configuration.
- * Modify the Percona XtraDB Cluster size parameter to add or remove Percona XtraDB Cluster members
- * Integrate with Percona Monitoring and Management (PMM) to seamlessly monitor your Percona XtraDB Cluster
- * Automate backups or perform on-demand backups as needed with support for performing an automatic restore
- * Supports using Cloud storage with S3-compatible APIs for backups
- * Automate the recovery from failure of a single Percona XtraDB Cluster node
- * TLS is enabled by default for replication and client traffic using Cert-Manager
- * Access private registries to enhance security
- * Supports advanced Kubernetes features such as pod disruption budgets, node selector, constraints, tolerations, priority classes, and affinity/anti-affinity
- * You can use either PersistentVolumeClaims or local storage with hostPath to store your database
- * Customize your MySQL configuration using ConfigMap.

16.13.1 Installation

Installation is performed by following the documentation installation instructions for [Kubernetes](#) and [OpenShift](#).

Last update: 2022-09-22