



VoltDB Kubernetes Administrator's Guide

Abstract

This book explains how to create and manage VoltDB database clusters using Kubernetes.

VoltDB Kubernetes Administrator's Guide

V15

Copyright © 2020-2025 Volt Active Data, Inc.

This document is published under copyright by Volt Active Data, Inc. All Rights Reserved.

The software described in this document is furnished under a license by Volt Active Data, Inc. Your rights to access and use VoltDB features are defined by the license you received when you acquired the software.

The VoltDB client libraries, for accessing VoltDB databases programmatically, are licensed separately under the MIT license.

Volt Active Data, VoltDB, and Active(N) are registered trademarks of Volt Active Data, Inc.

VoltDB software is protected by U.S. Patent Nos. 9,600,514, 9,639,571, 10,067,999, 10,176,240, and 10,268,707. Other patents pending.

This document was generated on September 30, 2026.

Table of Contents

Preface	viii
1. Structure of This Book	viii
2. Related Documents	viii
1. Introduction	1
1.1. Overview: Running VoltDB in Kubernetes	1
1.2. Setting Up Your Kubernetes Environment	3
1.2.1. Product Requirements	3
1.2.2. Configuring the Host Environment and Accounts	3
1.2.3. Configuring the Client	4
1.2.4. Granting Kubernetes Access to the Docker Repository	4
2. Configuring the VoltDB Database Cluster	6
2.1. Using Helm Properties to Configure Your Database	7
2.2. Configuring the Cluster	8
2.3. Configuring the Network Protocols	9
2.3.1. Configuring Network Protocols for XDCR	10
2.3.2. Configuring Multiple Network Interfaces	10
2.4. Configuring the Database	12
2.4.1. Configuring High Availability (K-Safety and Placement Groups)	12
2.4.2. Configuring Command Logging	15
2.4.3. Configuring Export	15
2.5. Configuring Logging	16
3. Starting and Stopping the Database	17
3.1. Starting the Cluster for the First Time	17
3.2. Stopping and Restarting the Cluster	17
3.3. Resizing the Cluster with Elastic Scaling	18
3.3.1. Increasing the Size of the Cluster	18
3.3.2. Decreasing the Size of the Cluster	18
3.3.3. Autoscaling the Cluster	18
3.4. Pausing and Resuming the Cluster	23
3.5. Starting More than One Cluster Within a Namespace	24
3.6. Stopping, Restarting, and Shutting Down Multiple Clusters Within a Namespace	25
4. Managing VoltDB Databases in Kubernetes	26
4.1. Managing the Cluster Using kubectl and helm	26
4.2. Managing the Database Using voltadmin and sqlcmd	27
4.2.1. Accessing the Database Interactively	27
4.2.2. Accessing the Database Programmatically	29
5. Updates and Upgrades	30
5.1. Updating the Database Schema	30
5.2. Updating the Database Configuration	31
5.2.1. Changing Database Properties on the Running Database	31
5.2.2. Changing Database Properties That Require a Restart	32
5.2.3. Changing Cluster Properties	33
5.3. Upgrading the VoltDB Software and Helm Charts	33
5.3.1. Updating Your Helm Repository	34
5.3.2. Updating the Custom Resource Definition (CRD)	34
5.3.3. Upgrading the VoltDB Operator and Software	34
5.3.4. Using In-Service Upgrade to Update the VoltDB Software	35
5.3.5. Updating VoltDB for XDCR Clusters	38
6. Monitoring VoltDB Databases in Kubernetes	39
6.1. Using Prometheus to Monitor VoltDB	39
7. Configuring Security in Kubernetes	41

7.1. Configuring User Accounts and Roles Within The Database	41
7.1.1. Assigning Administrative Access to the Volt Operator	41
7.1.2. Using Kubernetes Secrets to Store User Definitions	41
7.1.3. Updating User Account Secrets	42
7.2. Configuring TLS/SSL	43
7.2.1. Configuring TLS/SSL With the Helm Install --set-file Argument	44
7.2.2. Using Kubernetes Secrets to Store and Reuse TLS/SSL Information	44
7.2.3. Using Kubernetes cert-manager to Store TLS/SSL Certificates	45
7.2.4. Using Certificate Revocation Lists (CRLs)	46
7.2.5. Configuring TLS/SSL for the Volt Operator and VMC	47
7.3. Updating TLS/SSL Security Certificates	48
8. Cross Datacenter Replication in Kubernetes	49
8.1. Requirements for XDCR in Kubernetes	49
8.2. Choosing How to Establish a Network Mesh	49
8.3. Common XDCR Properties	51
8.4. Configuring XDCR in Local Namespaces	51
8.5. Configuring XDCR Using Load Balancers	52
8.5.1. Separate Load Balancers For Each Node (cluster.serviceSpec.perpod)	53
8.5.2. Single Load Balancer For Discovery with Virtual Networking Peering (cluster.serviceSpec.dr)	54
8.6. Configuring XDCR Using Node Ports for Replication	54
8.7. Configuring XDCR Using Network Services	55
9. Managing XDCR Clusters in Kubernetes	57
9.1. Removing a Cluster Temporarily	57
9.2. Removing a Cluster Permanently	57
9.3. Resetting XDCR When a Cluster Leaves Unexpectedly	58
9.4. Rejoining an XDCR Cluster That Was Previously Removed	58
A. Volt Diagnostics Tools	60
A.1. Starting and Stopping the Diagnostics Pod	60
A.1.1. Using the Diagnostics Pod With a Persistent Volume (PV)	61
A.1.2. Using the Diagnostic Tools Securely with TLS/SSL	61
A.2. Running the Diagnostic Tools	61
A.3. Diagnostic Tools	62
A.3.1. Collect (kcollect)	62
A.3.2. Mesh Monitor (meshmonitor, V1.5.0 and later)	62
A.3.3. Snapshot (ksnapshot, V1.3.0 and later)	63
A.3.4. Sqlcmd (ksqlcmd)	63
B. VoltDB Helm Properties	65
B.1. How to Use the Properties	65
B.2. Top-Level Kubernetes Options	66
B.3. Kubernetes Cluster Startup Options	66
B.4. Network Options	71
B.5. VoltDB Database Startup Options	72
B.6. VoltDB Database Configuration Options	73
B.7. Operator Configuration Options	80
B.8. Metrics Configuration Options	81
B.9. Volt Management Center (VMC) Configuration Options	81

List of Figures

1.1. Kubernetes/VoltDB Architecture	2
---	---

List of Tables

B.1. Top-Level Options	66
B.2. Options Starting with cluster.clusterSpec... ..	66
B.3. Options Starting with cluster.serviceSpec... ..	71
B.4. Options Starting with cluster.config... ..	73
B.5. Options Starting with cluster.config.deployment... ..	73
B.6. Options Starting with operator... ..	80
B.7. Options Starting with vmc... ..	81

List of Examples

5.1. Process for Upgrading the VoltDB Software	35
--	----

Preface

This book describes using Kubernetes and associated products to create and manage VoltDB databases and the clusters that host them. It is intended for database administrators and operators responsible for the ongoing management and maintenance of database infrastructure in a containerized environment.

This book is *not* a tutorial on Kubernetes or VoltDB. Please see “Related Documents” below for documents that can help you familiarize yourself with these topics.

1. Structure of This Book

This book is divided into 9 chapters and 2 appendices:

- Chapter 1, *Introduction*
- Chapter 2, *Configuring the VoltDB Database Cluster*
- Chapter 3, *Starting and Stopping the Database*
- Chapter 4, *Managing VoltDB Databases in Kubernetes*
- Chapter 5, *Updates and Upgrades*
- Chapter 6, *Monitoring VoltDB Databases in Kubernetes*
- Chapter 7, *Configuring Security in Kubernetes*
- Chapter 8, *Cross Datacenter Replication in Kubernetes*
- Chapter 9, *Managing XDCR Clusters in Kubernetes*
- Appendix A, *Volt Diagnostics Tools*
- Appendix B, *VoltDB Helm Properties*

2. Related Documents

This book assumes a working knowledge of Kubernetes, VoltDB, and the other technologies used in a containerized environment (specifically Docker and Helm). For information on developing and managing VoltDB databases, please see the manuals *Using VoltDB* and *VoltDB Administrator's Guide*. For new users, see the *VoltDB Tutorial*. For introductory information on the other products, please see their respective websites for appropriate documentation:

- Docker
- Helm
- Kubernetes

Finally, this book and all other documentation associated with VoltDB can be found on the web at <https://docs.voltactivedata.com/v15docs/>.

Chapter 1. Introduction

Kubernetes is an environment for hosting virtualized applications and services run in containers. It is designed to automate the management of distributed applications, with a particular focus on microservices. VoltDB is not a microservice — there is coordination between the nodes of a VoltDB cluster that requires additional attention. So although it is possible to spin up a generic set of Kubernetes "pods" to run a VoltDB database, additional infrastructure is necessary to realize the full potential of Kubernetes and VoltDB working together.

VoltDB Enterprise Edition provides additional services to simplify, automate, and unlock the power of running VoltDB within Kubernetes environments. There are six key components to the VoltDB Kubernetes offering, three available as open-source applications for establishing the necessary hosting environment and three provided by VoltDB to Enterprise customers. The three open-source products required to run VoltDB in a Kubernetes environment are:

- **Kubernetes** itself
- **Docker**, for managing the container images
- **Helm**, for automating the creation and administration of VoltDB in Kubernetes

In addition to these base requirements, VoltDB provides the following three custom components:

- **Pre-packaged docker image** for running VoltDB cluster nodes
- **The VoltDB Operator**, a separate utility (and docker image) for orchestrating the startup and management of VoltDB clusters in Kubernetes
- **Helm charts** for initializing and communicating with Kubernetes, the VoltDB Operator and its associated VoltDB cluster

The remainder of this chapter provides an overview of how these components work together to support running virtualized VoltDB clusters in a Kubernetes environment, the requirements for the host and client systems, and instructions for preparing the host environment prior to running VoltDB. Subsequent chapters provide details on configuring and starting your VoltDB cluster as well as common administrative tasks such as:

- Managing the running database with Helm and kubectl
- Updating the database schema, configuration, or the VoltDB software
- Configuring and managing security options for the database and auxiliary services
- Configuring and starting multiple clusters using cross datacenter replication (XDCR)


Finally, an appendix provides a full list of the Helm properties for configuring and controlling your VoltDB clusters.

1.1. Overview: Running VoltDB in Kubernetes

Kubernetes lets you create clusters of virtual machines, on which you run "pods". Each pod acts as a separate virtualized system or container. The containers are pre-defined collections of system and application components needed to run an application or service. Kubernetes provides the virtual machines, Docker

defines the containers, and Kubernetes takes responsibility for starting and stopping the appropriate number of pods that your application needs.

So the basic architecture for running VoltDB is a VoltDB database running on multiple instances of a Docker container inside a Kubernetes cluster. VoltDB also starts one or more auxiliary services as separate pods, such as the Volt Management Center.




However, out of the box, VoltDB and Kubernetes do not "talk together" and so there is no agreement on when pods are started and stopped and whether a VoltDB node is active or not. To solve this problem, VoltDB provides an additional service, the VoltDB Operator that manages the interactions between the VoltDB cluster, its auxiliary services, and the Kubernetes infrastructure. The Operator takes responsibility for initializing and starting the VoltDB server instances as appropriate, monitoring their health, and coordinating changes to the configuration.

To further simplify the process, VoltDB uses the open-source management product Helm to integrate Kubernetes, Docker, and VoltDB under a single interface. Helm uses "charts" to define complex management operations, such as configuring and starting the Kubernetes pods with the appropriate Docker images and then initializing and starting VoltDB on those pods. Simply by "installing" the appropriate Helm chart you can instantiate and run a VoltDB database cluster within Kubernetes using a single command.

Once the database is running, you can use standard VoltDB command line utilities to interact with and manage the database contents, such as modifying the schema or initiating manual snapshots. However, you will continue to use Helm to manage the server process and cluster on which the database runs, for activities such as stopping and starting the database. Figure 1.1, "Kubernetes/VoltDB Architecture" shows the overall architecture of using VoltDB, the VoltDB Operator, and Helm to automate running a VoltDB database within Kubernetes.

Figure 1.1. Kubernetes/VoltDB Architecture



1.2. Setting Up Your Kubernetes Environment

Before you can run VoltDB in a containerized environment, you must be sure your host systems and client are configured with the right software and permissions to support VoltDB. The following sections outline:

- What products are required on both the host environment and the local client you use to control Kubernetes and VoltDB
- How to configure the host environment and user accounts to run the VoltDB components
- How to configure your local client to control Kubernetes and the Helm charts
- How to set permissions in Kubernetes and Docker to allow access to the VoltDB components

1.2.1. Product Requirements

Before you start, you must make sure you have the correct software products and versions installed on both the host system and your local client. The host environment is the set of servers where Kubernetes is installed, whether they are systems you set up yourself or hosted by a third-party cloud service, such as the Google Cloud Platform or Microsoft Azure. The local client environment is the system, such as a desktop or laptop, you use to access the services.

The following are the software requirements for running VoltDB in Kubernetes.

Host Environment

- Kubernetes V1.28.x or later¹

VoltDB on Kubernetes has been validated for the following cloud service providers:

- AWS
- Azure
- Google Cloud
- OpenShift
- Tanzu

Client Environment

- Kubectl V1.27 or later²
- Helm V3.11.x or later

Optionally, you may want to install VoltDB on the client so you can use the **voltadmin** and **sqlcmd** command utilities to access the database remotely. If not, you can still use **kubectl** to create an interactive shell process on one of the server instances and run the utilities directly on the Kubernetes pods.

1.2.2. Configuring the Host Environment and Accounts

Once you have the necessary software installed, you must prepare the host environment to run VoltDB. This includes adding the appropriate Docker and chart repositories to Helm and configuring your host account with the permissions necessary to access those repositories.

¹Volt Active Data has been tested and verified against Kubernetes versions up to and including 1.35.x and is expected to operate correctly with all subsequent releases.

²Kubectl on the client must be within one minor version of Kubernetes in the host environment. For example, if Kubernetes is at version 1.31, Kubectl can be 1.30, 1.31, or 1.32. See the Kubernetes version skew documentation for further information.

First, you need accounts on the Kubernetes host environment and on the docker repository where the VoltDB images are stored, <https://docker.io>. To run the VoltDB Helm charts, your accounts must be set up with the following permissions:

- **Your Kubernetes host account** must have sufficient permissions to allocate persistent volumes and claims and create and manage pods.
- **Your Docker repository account** must have permission to access the VoltDB docker images. Access to the VoltDB docker images is assigned to VoltDB Enterprise customers on a per account basis. Contact VoltDB support for more information.

1.2.3. Configuring the Client

Next you must configure your client environment so you can communicate with and control Kubernetes and the Helm charts. First, install the Kubernetes and Helm command line interfaces, **kubectl** and **helm**. Next, configure the services to access the appropriate remote accounts and repositories.

The primary setup task for kubectl is creating the appropriate context for accessing the Kubernetes host you will be using. This is usually done as part of the installation or with a Kubconfig file and the **kubectl config** command. Once you have a context defined, you can use the **kubectl cluster-info** command to verify that your client is configured correctly.

For helm, you must add a link to the VoltDB chart repository, using the **helm repo add** command:

```
$ helm repo add voltdb \
    https://voltdb-kubernetes-charts.storage.googleapis.com
```

The first argument to the command ("voltdb") is a short name for referencing the repository in future commands. You can specify whatever name you like. The second argument is the location of the repository itself and must be entered as shown above.

Note

Helm first looks in local folders for charts you specify, then in the repositories. So if the short name you use matches a local directory, they can conflict and cause errors. In that case, you may want to choose a different name, such as "voltkube", to avoid any ambiguity. Then the chart locations you use in Helm commands would be "voltkube/voltdb" rather than "voltdb/voltdb" as shown in the examples.

1.2.4. Granting Kubernetes Access to the Docker Repository

Finally, you need to tell Kubernetes to access the Docker repository using the credentials for your Docker account. There are several ways to do this. You can specify your credentials on the helm command line each time you install a new VoltDB cluster. You can save the credentials in a YAML file with other parameters you pass to helm. Or you can set the credentials in a Kubernetes secret using kubectl.

The advantage of using a secret to store the credentials is that you only need to define them once and they are not easily discovered by others, since they are encrypted. To create a Kubernetes secret you use the **kubectl create secret** command, specifying the type of secret (*docker-registry*) and the name of the secret (which must be *dockerio-registry*), plus the individual credential elements as arguments:

```
$ kubectl create secret docker-registry dockerio-registry \
  --docker-username=johndoe \
  --docker-password='ThisIsASecret' \
```

```
--docker-email="jdoe@anybody.org"
```

Once you add the secret, you do not need to specify them again. If, on the other hand, you prefer to specify the credentials when you issue the helm commands to initialize the VoltDB cluster, you can supply them as the following helm properties using the methods described in Chapter 2, *Configuring the VoltDB Database Cluster*:

- `global.image.credentials.username`
- `global.image.credentials.password`

Chapter 2. Configuring the VoltDB Database Cluster

The two major differences between creating a VoltDB database cluster in Kubernetes and starting a cluster using traditional servers are:

- In Kubernetes, there is a single Helm command (`install`) that performs both the initialization and the startup of the database.
- You specify all the database and environment configuration with Helm properties rather than as separate YAML properties, environment variables, or command line arguments.

In fact, all of the configuration — including the configuration of the virtual servers (or pods), the server processes, and the database — is accomplished using Helm properties. Helm simplifies the process by coordinating all the different components involved, including Kubernetes, Docker, and VoltDB. By using the provided Helm charts, it is possible to start a default VoltDB cluster with a single command:

```
$ helm install mydb voltdb/voltdb \
  --set global.voltdbVersion=13.2.0 \
  --set-file cluster.config.licenseXMLFile=license.xml
```

For recent versions¹ of Volt Active Data, you can start the cluster with just four arguments to the `helm install` command:

1. **The release name** — The release name (*mydb* in the example) identifies the cluster and is used as a prefix for the corresponding Kubernetes' artifacts (such as pod and service names) as well as in Helm commands when managing the cluster. You can use any name you like to identify the cluster.
2. **The name of the Helm chart** — in the case of VoltDB, the chart name is the name you gave when you added the Volt chart repository to Helm followed by the chart name *voltdb*. Assuming you named the repository *voltdb*, the chart name is *voltdb/voltdb*.
3. **The VoltDB version** — You must also specify the version of Volt Active Data you want to use. This is the server software version.
4. **Your Volt Active Data license** — Finally, you must provide the Volt Active Data license file you received when you purchased Volt.

You specify the release and chart names as parameters to the command and you specify the VoltDB version and license file as values using command arguments such as `--set` and `--set-file`. All other properties — most notably the configuration properties — have default values that are used if not explicitly changed. For example, the default node (or replica) count for the cluster is three.

Note

Although the `global.voltdbVersion` property is required, it is a Helm property and can be set like any other, individually on the command line or in a properties file with other database and cluster settings. For the purposes of demonstration, many of the examples in this book, assume that the software version is set in a properties file as part of the overall configuration

¹Recent Volt versions (13.0.3, 12.3.4, 11.4.13, and 10.2.21 or later) use the simplified command for starting the database. For earlier versions, you must specify the specific chart and software versions to use. See the Volt Operator Release Notes for details.

However, a default cluster of three nodes and no schema or configuration is not particularly useful. So VoltDB provides Helm properties to let you customize every aspect of the database and cluster configuration, including:

- Cluster configuration, including size of the cluster, available resources, and so on
- Network configuration, including the assignment of ports and external mappings
- Database initialization options, including administration username and password, schema, and class files
- Database configuration, including the settings normally found in the YAML configuration files on non-Kubernetes installations

The following sections explain how to specify Helm properties in a properties file or on the command line, as well how to use those properties to make some of the most common customizations to your database. Later chapters explain how to configure specific features (such as security and XDCR). Appendix B, *VoltDB Helm Properties* provides a full list of the properties, including a brief description and the default value for each.

2.1. Using Helm Properties to Configure Your Database

First, it is useful to understand the different ways you can specify properties on the Helm command line. The following discussion is not intended as a complete description of Helm; only a summary to give you an idea of what they do and when to use them.

Helm offers three different ways to specify properties:

`--set`

The `--set` flag lets you specify individual property values on the command line. You can use `--set` multiple times or separate multiple property/value pairs with commas. For example, the following two commands are equivalent:

```
$ helm install mydb voltdb/voltdb \
  --set cluster.serviceSpec.clientPort=22222 \
  --set cluster.serviceSpec.adminPort=33333
$ helm install mydb voltdb/voltdb \
  --set cluster.serviceSpec.clientPort=22222,\
    cluster.serviceSpec.adminPort=33333
```

The `--set` flag is useful for setting a few parameters that change frequently or for overriding parameters set earlier in the command line (such as in a YAML file).

`--set-file`

The `--set-file` flag lets you specify the contents of a file as the value for a property. For example, the following command sets the contents of the file `license.xml` as the license for starting the VoltDB cluster:

```
$ helm install mydb voltdb/voltdb \
  --set-file cluster.config.licenseXMLFile=license.xml
```

As with `--set`, You can use `--set-file` multiple times or separate multiple property/file pairs with commas. The `--set-file` flag is useful for setting parameters where the value is too complicated to set directly on the command line. For example, the contents of the VoltDB license file.

--values, -f

The --values flag lets you specify a file that contains multiple property definitions in YAML format. Whereas properties set on the command line with --set use dot notation to separate the property hierarchy, YAML puts each level of the hierarchy on a separate line, with indentation and followed by a colon. For example, the following YAML file and --values flag set the same two properties shown in the --set example above:

```
$ cat ports.yaml
cluster:
  serviceSpec:
    clientPort: 22222
    adminPort: 33333
$ helm install mydb voltdb/voltdb \
  --values ports.yaml
```

YAML files are extremely useful for setting multiple properties with values that do not change frequently. You can also use them to group properties (such as port settings or security) that work together to configure aspects of the database environment.

You can use any of the preceding techniques for specifying properties for the VoltDB Helm charts. In fact, you can use each method multiple times on the command line and mixed in any order. For example, the following example uses --values to set the database configuration and ports, --set -file to identify the license, and --set to specify the number of nodes requested:

```
$ helm install mydb voltdb/voltdb \
  --values dbconf.yaml,dbports.yaml \
  --set-file cluster.config.licenseXMLFile=license.xml \
  --set cluster.clusterSpec.replicas=5
```

There are a few important points about how Helm and the Volt Operator interpret YAML you should pay attention to when configuring your database clusters. Specifically:

- Using Helm version 3.x, large numbers (six digits or more) could be converted to scientific E notation (for example from 4000000 to 4.0e+06), which is not recognized by the Volt server. To avoid your input being reformatted, enclose large numbers in quotation marks so they are interpreted as a string value, which is passed as-is to the Volt server. For Helm V4.0 and later, this workaround is not necessary.
- In YAML it is possible to set specific elements of a list by using array-like notation. For example, the following YAML identifies the third element of the list of export connectors: *cluster.config.deployments.export.configurations[2]*. However, in several cases, the Helm YAML syntax includes subproperties of list elements. If you use the **helm upgrade** command to reset one specific subelement of a list, it resets the entire list element and you need to include any other subproperties that are not currently set by default. In general, it is easier and less prone to error if you re-specify the entire list and its elements, ensuring that all subproperties are set as desired.

2.2. Configuring the Cluster

Many of the configuration options that are performed through hardware configuration, system commands or environment variables on traditional server platforms are now available through Helm properties. Most of these settings are listed in Section B.3, “Kubernetes Cluster Startup Options”.

Hardware Settings

Hardware settings, such as the number of processors and memory size, are defined as Kubernetes image resources through the Helm `cluster.clusterSpec.resources` property. Under re-

`sources`, you can specify any of the YAML properties Kubernetes expects when configuring pods within a container. For example:

```
cluster:
  clusterSpec:
    resources:
      requests:
        cpu: 500m
        memory: 1000Mi
      limits:
        cpu: 500m
        memory: 1000Mi
```

System Settings

System settings that control process limits that are normally defined through environment variables can be set with the `cluster.clusterSpec.env` properties. For example, the following YAML increases the Java maximum heap size and disables the collection of JVM statistics:

```
cluster:
  clusterSpec:
    env:
      VOLTDB_HEAPMAX: 3072
      VOLTDB_OPTS: -XX:+PerfDisableSharedMem
```

Note

There are no Volt database or cluster properties specific to the server hardware architecture. However, when running Volt Active Data on ARM64 processors, your Kubernetes host provider may define certain Kubernetes attributes, such as taints, specific to ARM for which your Helm chart must make allowances. In this case, you can use the `tolerations` properties associated with the individual Volt services — including `cluster.clusterSpec.tolerations`, `operator.tolerations`, and `vmc.tolerations` — to provide the necessary accommodations.

One system setting that is *not* configurable through Kubernetes or Helm is whether the base platform has Transparent Huge Pages (THP) enabled or not. This is dependent of the memory management settings on the actual base hardware on which Kubernetes is hosted. Having THP enabled can cause problems with memory-intensive applications like VoltDB and it is strongly recommended that THP be disabled before starting your cluster. (See the section on Transparent Huge Pages in the *VoltDB Administrator's Guide* for an explanation of why this is an issue.)

If you are not managing the Kubernetes environment yourself or cannot get your provider to modify their environment, you will need to override VoltDB's warning about THP on startup by setting the `cluster.clusterSpec.additionalArgs` property to include the VoltDB start argument to disable the check for THP. For example:

```
cluster:
  clusterSpec:
    additionalStartArgs:
      - "--ignore=thp"
```

2.3. Configuring the Network Protocols

Kubernetes has the ability to configure the networks serving your pods using IPv4, IPv6, or both (what is known as a *dual stack*). How the network protocols are configured differs depending on the flavor of

Kubernetes and the hosting service you use. But once established, Volt can use whichever protocol is available. In the case of a dual stack, Volt will choose one or the other protocol for the individual interfaces as it needs.

2.3.1. Configuring Network Protocols for XDCR

Since XDCR involves communication between clusters, it is important you know which protocol is in use so you can configure the XDCR connections. So the Volt Helm chart provides the `.ipFamilies` property so you can choose which protocol to use for XDCR communication. There are two flavors of the `.ipFamilies` property, depending on how you are configuring XDCR. If you are establishing an XDCR environment within a single namespace (as described in Section 8.4, “Configuring XDCR in Local Namespaces”), you set the IP version in the `cluster.serviceSpec.dr.ipFamilies` property, as in the following example that sets the IP family to IPv6:

```
cluster:
  serviceSpec:
    dr:
      enabled: true
      ipFamilies:
        - IPv6
```

If you are configuring per pod XDCR, you select the IP version in the `cluster.serviceSpec.perPod.dr.ipFamilies` property; as in this example:

```
cluster:
  serviceSpec:
    perPod:
      dr:
        enabled: true
        ipFamilies:
          - IPv6
```

It is possible to customize the network environment even further using third-party networking services, such as Multus — which is a container networking service (CNI) plugin that gives pods access to multiple networks. Again, how you configure third-party services is beyond the scope of this manual and differs from one service to the next. However, these services provide instructions to the pods through annotations that you include in your Helm chart. Specifically, Volt supports annotation properties that affect the XDCR settings, which Multus can use to specify alternate networks, as in the following example:

```
cluster:
  serviceSpec:
    dr:
      annotations:
        aviinfra.setting.ako.vmware.com/name: my-infra-sets
  clusterSpec:
    additionalAnnotations:
      k8s.v1.cni.cncf.io/networks: my-custom-network
```

See the Multus documentation for more information.

2.3.2. Configuring Multiple Network Interfaces

The annotations shown in the preceding section tell a container networking service (CNI) such as Multus which secondary networks to attach to each pod. When a pod has more than one network interface, VoltDB

can decide which interfaces to assign to what roles. By default, VoltDB advertises its primary address (eth0) for both clients and remote XDCR clusters. If you want client or replication traffic to use a secondary network instead, assign the specific interface to the corresponding `cluster.clusterSpec.network` property.

You configure each type of connection separately. The `network.client` property sets the address that topology-aware client applications connect to, and `network.dr` sets the address a remote cluster uses to establish cross datacenter replication (XDCR). You can configure either or both, and the two need not use the same interface.

For each connection, the `resolveBy` property specifies how VoltDB selects the address from the pod's available interfaces when it starts. Setting `resolveBy` to `cidr` selects the interface whose address falls within the subnet you specify with the `cidr` property; setting `resolveBy` to `interface` selects the address on the interface you name with the `interface` property. (Multus typically names secondary interfaces `net1`, `net2`, and so on.) For example, the following configuration assigns the client address from a secondary subnet:

```
cluster:
  clusterSpec:
    network:
      client:
        resolveBy: cidr
        cidr: 192.168.100.0/24
```

By default the network property changes only the address VoltDB advertises; the database still listens on all interfaces. To also restrict a connection's listener to its resolved interface — so that its traffic cannot travel over the primary network — set the `bind` property to `true` on that YAML block. And on a dual-stack pod, where an interface can have both an IPv4 and an IPv6 address, set `ipFamily` to either `ipv4` (the default) or `ipv6`, to choose which one VoltDB advertises. This per-connection setting is distinct from the `cluster.serviceSpec... ipFamilies` property described in the previous section, which selects the protocol for the XDCR service rather than the address resolved on the pod.

The following example selects each network by name, binds each listener to its interface, and advertises the IPv6 address:

```
cluster:
  clusterSpec:
    network:
      client:
        resolveBy: interface
        interface: net1
        bind: true
        ipFamily: ipv6
      dr:
        resolveBy: interface
        interface: net2
        bind: true
        ipFamily: ipv6
```

Finally, the `timeout` property specifies how long startup waits for the chosen address to appear. (The timeout is specified in seconds with a default of 120, or 2 minutes.) If VoltDB cannot resolve the address in the allotted time — for example, because of an incorrect subnet or interface name, or a secondary network that is not attached — it stops attempting to resolve the address. The node does not start and the pod stays in the *Running* state without becoming *Ready*. You can use the pod's log to see whether and how each address was resolved.

2.4. Configuring the Database

In addition to configuring the environment VoltDB runs in, there are many different characteristics of the database itself you can control. These include mapping network interfaces and ports, selecting and configuring database features, and identifying the database schema, class files, and security settings.

The network settings are defined through the `cluster.serviceSpec` properties, where you can choose the individual ports and choose whether to expose them through the networking service. For example, the following YAML file disables exposure of the admin port and assigns the externalized client port to 31313:

```
cluster:
  serviceSpec:
    type: NodePort
    adminPortEnabled: false
    clientPortEnabled: true
    clientNodePort: 31313
```

The majority of the database configuration options for VoltDB are declared using YAML for both bare metal and Kubernetes. In general, almost all of the options are identical on both, except that on bare metal the properties start with the prefix `deployment`, whereas Helm properties start with `cluster.config.deployment`. So, for example, where the number of sites per host is defined on bare metal as:

```
deployment:
  cluster:
    sitesperhost: {n}
```

It is defined in Kubernetes as:

```
cluster:
  config:
    deployment:
      cluster:
        sitesperhost: {n}
```

The following sections give examples of defining common database configurations options using YAML. See Section B.6, “VoltDB Database Configuration Options” for a complete list of the Helm properties available for configuring the database.

2.4.1. Configuring High Availability (K-Safety and Placement Groups)

Volt Active Data provides high availability through *K-safety*, where copies of each partition are distributed to different nodes in the database cluster. If a node fails, the database can continue to operate because there are still copies of every partition within the cluster. The amount of durability depends on the K factor. So a K factor of one means that the cluster is guaranteed to survive one node (or pod) failing, a factor of two guarantees two nodes, and so on. (See the chapter on Availability in the *Using VoltDB* manual for more information on how K-safety works.)

You set the K-safety factor using the `cluster.config.deployment.cluster.kfactor` property when configuring your database. For example, the following YAML sets the K-safety factor to two:

```
cluster:
```

```
clusterSpec:
  replicas: 6
config:
  deployment:
    cluster:
      sitesperhost: 8
      kfactor: 2
```

Note that the number of replicas must be at least as large as the K factor plus one (K+1) and K-safety is most effective if the number of replicas times the number of sites per host is a multiple of K+1.

The combination of K-safety and Kubernetes provides an automated, self-healing system where K-safety ensures the cluster survives individual nodes failing and Kubernetes manages the automated recreation of the pods when they fail so the database can be restored to a full complement of nodes as soon as possible. However, to take full advantage of this capability you need to ensure the Kubernetes infrastructure is configured correctly to distribute the Volt servers evenly and that Volt uses information about the configuration to manage the distribution of partitions within the database. The following sections explain how to use Kubernetes configuration options, such as affinity and spread constraints, and Volt placement groups to achieve maximum availability.

2.4.1.1. Configuring Kubernetes Clusters for High Availability (Spread Constraints and Affinity)

K-safety ensures the database cluster can survive at least a certain number of node failures. However, to reduce the risk of larger scale outages, you need to make sure that the Volt servers are distributed in such a way to minimize the impact of external outages. In particular, you want to ensure that each Volt server pod runs on a separate Kubernetes node (so that a Kubernetes node failure cannot impact multiple pods) and that the pods are, as much as possible, evenly distributed among the availability zones in use.

By default, the Volt Operator establishes Kubernetes affinity and anti-affinity rules such that no two Volt server pods can run on the same Kubernetes node. So, normally, you do not need to take any actions to make this happen. However, if you are overriding the Operator's default configurations, you will need to make sure your custom Kubernetes configuration includes this behavior.

When using multiple availability zones, you should also adjust the Kubernetes configuration — specifically the spread constraints — so that the Volt server pods are evenly distributed among the zones. This makes it possible to avoid the database failing due to the loss of any one zone that contains an unbalanced and excessive number of Volt server processes. You can define the distribution of server pods within your Helm configuration using the `cluster.clusterSpec.topologySpreadConstraints` property. The following example demonstrates how to do this, using the label selector to identify the Volt server processes.

```
cluster:
  clusterSpec:
    topologySpreadConstraints:
      - topologyKey: topology.kubernetes.io/zone
        whenUnsatisfiable: DoNotSchedule
        maxSkew: 1
        labelSelector:
          matchLabels:
            name: voltdb-cluster
```

If you are running multiple databases within a single namespace, you should consider replacing the last line of the configuration, "name: voltdb-cluster", with an identifier that is specific to the cluster being

configured. For example, if the cluster release name is *mydb*, the last line of the configuration should read "voltdb-cluster-name: mydb-voltdb-cluster".

2.4.1.2. Cloud Native Placement Groups

K-safety guarantees the minimum number of nodes that can fail without stopping the database. Configuring Kubernetes affinity and spread constraints to evenly distribute the database server pods reduces the overall threat of external failures taking down the database. However, to fully maximize the availability, Volt needs to use knowledge about the Kubernetes configuration to intelligently distribute the individual copies of the partitions among those servers.

The cluster may survive more failures than just the minimum guaranteed by K-safety depending on how the partitions are distributed and which nodes fail. *Placement groups* are a mechanism for providing more context concerning the hardware environment to improve the likelihood of the cluster surviving multiple failures. For example, if you tell Volt certain nodes are in the same region and zone (i.e. in the same placement group), it avoids placing all copies of any partition on those nodes, so if the zone fails, the database can survive.

Because you do not control exactly where each pod is created in Kubernetes, Volt can use its knowledge of the Kubernetes availability zones and regions² to automate the placement groups and minimize the potential of an infrastructure failure taking the database down with it. You enable cloud native placement groups in Kubernetes by setting the property `cluster.clusterSpec.useCloudNativePlacementGroup` to "true". For cloud native placement groups to be effective, the cluster configuration must meet the following requirements:

- The cluster must be distributed over three or more regions or availability zones.
- The number of nodes (or *replicas*) must be a multiple of the number of availability zones.
- The number of availability zones must be a multiple of K+1.

For example, the following configuration assumes the cluster is distributed across four availability zones:

```
cluster:
  clusterSpec:
    replicas: 8
    useCloudNativePlacementGroup: true
  config:
    deployment:
      cluster:
        sitesperhost: 8
        kfactor: 1
```

Once the database is running, you can use the `@Statistics` system procedure with the `HOST` selector to determine where each node is running and what partitions are running on that node. In addition, if one or more nodes go down, the "SAFETOSTOP" column lets you know which of the remaining nodes could safely be stopped without endangering the cluster as a whole.

```
$ sqlcmd
1> execute @Statistics HOST;
TIMESTAMP      HOST_ID HOSTNAME      PARTITIONS      LEADERS      PL
```

²Placement groups depend on the Kubernetes labels `topology.kubernetes.io/region` and `topology.kubernetes.io/zone`, which are defined automatically by most commercial cloud providers. If you are using a custom cloud deployment, you will need to make sure these labels are declared appropriately before enabling cloud native placement groups.

```

1677777171869      0 mydb-voltdb-cluster-0 24,25,26,27,28,29,30,31 25,27,29,31 ea
1677777171870      1 mydb-voltdb-cluster-1 8,9,10,11,12,13,14,15 8,10,12,14 ea
1677777171870      2 mydb-voltdb-cluster-2 8,9,10,11,12,13,14,15 9,11,13,15 ea
1677777171870      3 mydb-voltdb-cluster-3 16,17,18,19,20,21,22,23 16,18,20,22 ea
1677777171870      4 mydb-voltdb-cluster-4 16,17,18,19,20,21,22,23 17,19,21,23 ea
1677777171870      5 mydb-voltdb-cluster-5 24,25,26,27,28,29,30,31 24,26,28,30 ea
1677777171870      6 mydb-voltdb-cluster-6 0,1,2,3,4,5,6,7 0,2,4,6 ea
1677777171870      7 mydb-voltdb-cluster-7 0,1,2,3,4,5,6,7 1,3,5,7 ea
(Returned 8 rows in 0.01s)
TIMESTAMP      PLACEMENTGROUP SAFETOSTOP
-----
1677777171882 east--zone3      true
1677777171882 east--zone2      true
1677777171882 east--zone1      true
1677777171882 east--zone4      true

```

2.4.2. Configuring Command Logging

Command logging provides durability of the database content across failures. You can control the level of durability as well as the length of time required to recover the database by configuring the type of command logging and size of the logs themselves. In Kubernetes this is done with the `cluster.config.deployment.commandlog` properties. The following example enables synchronous command logging and sets the log size to 3,072 megabytes and the frequency to 1,000 transactions:

```

cluster:
  config:
    deployment:
      commandlog:
        enabled: true
        synchronous: true
        logsize: 3072
        frequency:
          transactions: 1000

```

2.4.3. Configuring Export

Export simplifies the integration of the VoltDB database with external databases and systems. You use the export configuration to define external "targets" the database can write to. In Kubernetes you define export targets using the `cluster.config.deployment.export.configurations` property. Note that the `configurations` property can accept multiple configuration definitions. In YAML, you specify a list by prefixing each list element with a hyphen, even if there is only one element. The following example defines one export target, *eventlog*, using the file export connector:

```

cluster:
  config:
    deployment:
      export:
        configurations:
          - target: eventlog
            type: file
            properties:
              type: csv
              nonce: eventlog

```

2.5. Configuring Logging

VoltDB uses Log4J for logging messages while the database is running. The chapter on "Logging and Analyzing Activity in a VoltDB Database" in the *VoltDB Administrator's Guide* describes some of the ways you can customize the logging to meet your needs, including changing the logging level or adding appenders. Logging is also available in the Kubernetes environment and is configured using a Log4J properties file. However, the default configuration and how you set the configuration when starting or updating the database in Kubernetes is different than as described in the *Administrator's Guide*.

Before you attempt to customize the logging, you should familiarize yourself with the default settings. The easiest way to do this is to extract a copy of the default configuration from the Docker image you will be using. The following commands create a docker container without actually starting the image, extract the configuration file to a local file (`my-k8s-log4j.properties` in the example), then delete the container.

```
$ ID=$(docker create voltdb/voltdb-enterprise)
$ docker cp ${ID}:/opt/voltdb/tools/kubernetes/voltserver-k8s-log.properties. \
  my-k8s-log4j.properties
$ docker rm $ID
```

Once you extract the default configuration and made the changes you want, you are ready to specify your new configuration on the Helm command to start the database. You do this by setting the `cluster.config.log4jcfgFile` property. For example:

```
$ helm install mydb voltdb/voltdb \
  --values myconfig.yaml \
  --set cluster.clusterSpec.replicas=5 \
  --set-file cluster.config.licenseXMLFile=license.xml \
  --set-file cluster.config.log4jcfgFile=my-k8s-log4j.properties
```

Similarly, you can update the logging configuration on a running cluster by using the `--set-file` argument on the Helm upgrade command:

```
$ helm upgrade mydb voltdb/voltdb --reuse-values \
  --set-file cluster.config.log4jcfgFile=my-k8s-log4j.properties
```

Chapter 3. Starting and Stopping the Database

The key to managing VoltDB clusters in Kubernetes is to let the Helm charts do the work for you. You can use **helm** commands to perform all basic activities for running a database. This chapter explains how to use helm commands to:

- Start the cluster for the first time
- Stop and restart the cluster
- Resize the cluster
- Pause and resume
- Start multiple clusters within one Kubernetes namespace

Subsequent chapters explain how to manage the database once it is running, how to modify the database and cluster configuration, and how to upgrade the VoltDB software itself.

3.1. Starting the Cluster for the First Time

As described in Chapter 2, *Configuring the VoltDB Database Cluster* you can customize every aspect of the database and the cluster using Helm properties and the configuration can be as simple or as complex as you choose. But once you have determined the configuration options you want to use, actually initializing and starting the database cluster is a single command, **helm install**. For example:

```
$ helm install mydb voltdb/voltdb \
  --values myconfig.yaml \
  --set global.voltdbVersion=13.2.0 \
  --set-file cluster.config.licenseXMLFile=license.xml \
  --set cluster.clusterSpec.replicas=5
```

When running VoltDB in Kubernetes, it is important to always use the latest versions of the Volt Operator and custom resource definition (CRD). The first time you start a database in Kubernetes, Helm automatically uses the latest operator and CRD. Helm continues to use the latest operator on subsequent installs. However, it does not automatically update the CRD; you must do this manually. Therefore, if you have used Helm to run VoltDB before, it is best to update the CRD before starting a new database, as described in Section 5.3.2, “Updating the Custom Resource Definition (CRD)”.

It is also possible to identify specific versions of the Volt Operator and/or chart when starting VoltDB. However, if you do, be sure to use matching versions of the Operator and chart to avoid any potential compatibility issues.

3.2. Stopping and Restarting the Cluster

Once the cluster is running (what Helm calls a “release”), you can adjust the cluster to stop it, restart it, or resize it, by “upgrading” the release chart, specifying the new value for the number of nodes you want. You upgrade the release using much the same command as you do to start it, except rather than repeating the configuration, you can use the **--reuse-values** flag. So, for example, to stop the cluster, you simply set the number of replicas to zero, reusing all other parameters:

```
$ helm upgrade mydb voltdb/voltdb      \  
  --reuse-values                        \  
  --set cluster.clusterSpec.replicas=0
```

To restart the cluster after you stop it, you reset the replica count to five, or whatever you set it to when you initially defined and started it:

```
$ helm upgrade mydb voltdb/voltdb      \  
  --reuse-values                        \  
  --set cluster.clusterSpec.replicas=5
```

3.3. Resizing the Cluster with Elastic Scaling

Elastic scaling lets you increase or decrease the size of your cluster on the fly, without requiring any down-time. You initiate elastic scaling by changing the value of the `cluster.clusterSpec.replicas` property.

For both increasing and decreasing the size of the cluster, the change in the number of nodes must meet the configuration requirements for the cluster's K factor. Specifically, you must add or remove K+1 nodes at a time.

3.3.1. Increasing the Size of the Cluster

To scale up the cluster you add nodes by upgrading the release, specifying the new number of nodes you want. Of course, the new value must meet the requirements for elastically expanding the cluster, as set out in the discussion of adding nodes to the cluster in the *VoltDB Administrator's Guide*. So, for example, to increase a five node cluster with a K-safety factor of one by two nodes, you can set the replica count to seven:

```
$ helm upgrade mydb voltdb/voltdb      \  
  --reuse-values                        \  
  --set cluster.clusterSpec.replicas=7
```

3.3.2. Decreasing the Size of the Cluster

To scale down or "shrink" the cluster, you upgrade the release specifying the new number of nodes you need. The new value must meet the requirements for K-safety. Specifically, you can only remove K+1 nodes at a time. So if your cluster has a K-safety factor of one, you must remove two nodes to reduce the size of the cluster. For example, if you have a five node cluster with K=1, you can shrink the cluster by setting the replica count to three:

```
$ helm upgrade mydb voltdb/voltdb      \  
  --reuse-values                        \  
  --set cluster.clusterSpec.replicas=3
```

If you want to reduce the cluster by more than K+1 nodes (for example, going from seven to three nodes in the preceding example), you must iterate the resizing operation in steps of K+1 nodes. In our example that means first reducing the seven node cluster to five nodes. Then once the first resizing operation is complete, you can perform a second resizing operation to reduce the cluster to three nodes.

3.3.3. Autoscaling the Cluster

In dynamic environments there can be significant variation in the workload, based on anything from time of day to external events that drive usage. For example, traffic monitoring applications see the majority of

their activity during rush hour with very little at other times of day. For applications where usage can spike and dip it is useful to be able to scale up and scale down the cluster to meet the needs of the workload. Automatically adjusting the cluster size based on workload or other system performance metrics is called *autoscaling*.

To set up autoscaling, you need to define thresholds on one or more metrics that are used to trigger the resizing. For example, you might want to resize the cluster based on capacity, the amount of memory in use — scaling up when memory usage exceeds a certain level and scaling down if memory drops below a minimum size. Alternatively, you could schedule autoscaling based on throughput or CPU usage.

When you select a unit to measure, you also define the thresholds at which the cluster should scale up or down for that metric. For example, if you are using memory consumption as the trigger, you set the thresholds as the resident set size (RSS) in bytes, so that the cluster will scale up if too much memory is being used and scale down if too little is used.

Having enabled autoscaling, if the cluster ever does exceed the upper limit, the Volt operator will automatically start an elastic resize operation to increase the size of the cluster by K+1 nodes. Similarly, if the memory usage drops below the lower limit, the operator will start an elastic downsizing of the cluster by K+1 nodes.

In addition to defining the trigger metrics and their associated upper and lower thresholds, you can control other aspects of the autoscaling process, such as:

- The minimum and maximum size for the cluster
- How frequently autoscaling status is reported in the logs
- How long the metric must remain beyond the threshold before actually starting to autoscale the cluster

The following sections explain how to configure autoscaling, how to monitor autoscaling while it is in process, and how to troubleshoot potential problems if autoscaling fails.

3.3.3.1. Enabling Autoscaling

You configure autoscaling using the `cluster.clusterSpec.autoScaling.*` properties. First, you must set the `cluster.clusterSpec.autoScaling.enabled` property to *true*. Next, you select one or more metrics to measure as triggers for autoscaling. At the same time, you define the upper and lower thresholds using the `.scaleUp` and `.scaleDown` properties. There are currently four metrics you can choose from:

- CPU Usage (`cpu`) — measured in percentage of total CPU currently in use
- Memory Usage — measured by the resident set size in megabytes (`rss`) or percentage of total memory (`memused`)
- Latency (`tps`) — measured in the average number of transactions per second
- Idle Time (`idletime`) — measured in the percentage of time the partitions were idle (that is, not processing any transactions because the queues were empty)

You must define at least one metric as a trigger, but you can define more than one if you like. For example, the following Helm chart enables autoscaling based on both TPS and CPU, using 50K TPS or 75% CPU usage as the upper threshold and 10K TPS or 25% CPU usage as the lower threshold:

```
cluster:
  clusterSpec:
    autoScaling:
```

```
enabled: true
metrics:
  tps:
    scaleUp: 50000
    scaleDown: 10000
  cpu:
    scaleUp: 75
    scaleDown: 25
```

3.3.3.1.1. Setting Appropriate Autoscaling Thresholds

Defining appropriate thresholds is critical to effective autoscaling. What metrics you choose and what limits to set depend on the needs of your specific application. However, you have to be careful that the thresholds are both not too far apart that your application hits a resource constraint before scaling can take effect and not so close together that autoscaling ends up bouncing between scale up and scale down.

On the one hand, autoscaling does not take effect immediately. The trigger metrics are monitored periodically, so there can be a delay between when a threshold is crossed and when the Volt operator detects the event. Then actually scaling the cluster takes time as well. If a metric is steadily increasing, you do not want it to reach its physical maximum before the additional nodes are operational. So, for example, it would be dangerous to set the `cpu.scaleUp` threshold at 90%, since your cluster could easily reach 100% before the autoscaling is complete. Be sure to leave enough headroom for further growth until the elastic scaling operation can run to completion.

On the other hand, if you set the limits too close together, there is the danger that scaling up the cluster will reduce the trigger metric to the point where it now drops below the scale down threshold. For example, if you set the TPS limits at 20K and 40K, scaling up a 3-node, $K=2$ cluster to 6 nodes could easily cut the TPS in half, risking it dropping below the scale down threshold of 20K. The result would be a cluster that is constantly switching between scaling up and scaling down, seriously impinging on the cluster's ability to process requests.

The one other unusual situation to consider, when monitoring multiple metrics, is if one metric crosses the threshold in one direction while a second metric crosses the threshold in the opposite direction. For example, using the Helm chart from the preceding section, if the TPS upper threshold of 50K is crossed while at the same time the CPU usage drops below the lower threshold of 25%, the Volt operator will not take any action. Autoscaling will not be triggered unless *only* the upper or lower thresholds are crossed. This rule also applies when using a stabilization window (as described in the next section); if two thresholds are crossed in opposite directions during the stabilization window, the resize operation is canceled and the stabilization window reset.

3.3.3.1.2. Controlling the Autoscaling Process

Finally, there are additional properties that let you control the behavior of autoscaling. These include:

- **`cluster.clusterSpec.autoScaling.minReplicas`** and **`cluster.clusterSpec.autoScaling.maxReplicas`**

These properties specify the minimum and maximum number of nodes in the cluster. Autoscaling will not resize the cluster beyond these values, even if the threshold of a trigger metric is exceeded. The default minimum is $K+1$. The default maximum number of nodes is 16.

- **`cluster.clusterSpec.autoScaling.stabilizationWindow`**

There are actually two properties (`.scaleUp` and `.scaleDown`) specified in seconds that define how long after a threshold is crossed before the Volt operator actually starts the elastic resizing. If the metric drops back across the threshold during that window, the elastic operation is canceled and the stabilization "clock" reset. Providing a stabilization window allows applications with more dynamic workloads

to cross the thresholds temporarily without triggering an elastic resize operation until the threshold is crossed for an extended period of time. The default stabilization window is 10 minutes (600 seconds).

- **cluster.clusterSpec.autoScaling.maxRetries** and **cluster.clusterSpec.autoScaling.retryTimeout**

Specify how many times an autoscaling resize operation is retried after it fails and how long to wait for such operations to start before deciding that it has failed. The default is not to retry any failed resizing (a value of zero) and the retry timeout is 60 seconds. See Section 3.3.3.3, “Troubleshooting Autoscaling and Recovering From Errors” for more information about troubleshooting autoscaling failures.

- **cluster.clusterSpec.autoScaling.notificationInterval**

This property specifies (in seconds) how frequently the Volt operator updates the logs and statistics during an autoscaling event. The default notification interval is zero, or no notifications.

3.3.3.2. Monitoring Autoscaling

Once autoscaling begins, it takes time for the cluster to either elastically expand or shrink, which can vary significantly depending on the size of the data, the current workload, and other circumstances. During this period, it is a good idea to monitor the process to ensure the elastic operations are proceeding as expected.

The operator periodically updates the status of the operation in the operator log, the Kubernetes cluster status, and as Kubernetes events. You can see the operator logs using the **kubectl logs** command. For example, if the Helm release name is *mydb*:

```
$ kubectl logs -f deploy/mydb-voltdb-operator
```

To see the Kubernetes events associated with the operator, you can use the **kubectl events** command:

```
$ kubectl events --for deploy/mydb-voltdb-operator
```

Or you can see both the cluster status and Kubernetes events using the **kubectl describe** command:

```
$ kubectl describe voltdbcluster mydb-voltdb-cluster
[ . . . ]
Status:
  Cluster State:
    Auto Scaling:
      Desired Replicas: 3           ❶
      Metrics:           ❷
      Tps:
      State: Monitoring auto-scaling metrics  ❸
```

You can use whichever method you find most useful. However, the following examples use the cluster status because it is the easiest to read. In the preceding example, where a three node, K=1 cluster has autoscaling enabled, the status display tells you:

1. The number replicas the operator expects. Before any autoscaling occurs, this matches the node count for the cluster (3).
2. Which metrics are being monitored. In this case, only the TPS metric is being used.
3. The overall state of autoscaling. In this case, autoscaling is enabled and the metrics are being monitored.

If a metric crosses an associated threshold, the status changes to indicate which metric it is and when the threshold was crossed:

Status:

```
Cluster State:
Auto Scaling:
  Desired Replicas: 3 ❶
  Metrics:
    Tps:
      Direction:      ScaleUp ❷
      Last Value:     22981
      Time Threshold Crossed: 2024-02-02T19:27:05.870Z
    State:            Monitoring auto-scaling metrics ❸
```

The number of replicas (1) remains the same, because the stabilization window is in effect. But the metrics section (2) now includes information on which threshold has been triggered and when. While the status line still reports it is monitoring the metrics to make sure it stays over the threshold for the duration of the stabilization window.

If the metrics stays above (or below) the threshold when the stabilization window expires, the autoscaling event begins and the status changes to reflect the new state:

```
Status:
Cluster State:
Auto Scaling:
  Desired Replicas: 5 ❶
  Direction:      ScaleUp
  Metrics:
    Tps:
      Direction:      ScaleUp ❷
      Last Value:     28617
      Time Threshold Crossed: 2024-02-02T19:38:03.763Z
    State:            Scaling cluster ❸
  Time Scaling Notified: 2024-02-02T19:38:34.695Z
  Time Scaling Started:  2024-02-02T19:38:34.695Z
```

The target number of replicas (1) increases by $K+1$, to 5. Note this is the desired number of nodes, *not* the current number. It also reports the direction the cluster is resizing (in this case, scaling up). The metrics information (2) continues to report the current value and when the threshold was crossed triggering the autoscale activity. Finally, the status line (3) now reports that the cluster is elastically resizing to expand the cluster.

The cluster status will continue to report this information throughout the elastic resizing, which can take minutes, or even hours, depending on how much data must be moved and how busy the database is processing transactions. But once the resize is complete, the status returns to its initial stage, except the number of desired replicas now matches the new cluster size, 5.

```
Status:
Cluster State:
Auto Scaling:
  Desired Replicas: 5
  Metrics:
    Tps:
  State: Monitoring auto-scaling metrics
```

3.3.3.3. Troubleshooting Autoscaling and Recovering From Errors

The `autoScaling.maxRetries` and `autoScaling.retryTimeout` properties give the Volt operator some flexibility in recovering from issues that might arise during an autoscaling event. However,

not all failures are transitory. For example, there may be provisioning or configuration issues that are stopping the cluster from expanding. Which it is why it is important to know what to do if autoscaling does not operate correctly.

The first step is to monitor the autoscaling activity — either proactively or by reviewing the logs and status after the fact. If the maximum number of retries is exceeded or a scaling operation fails to restart, the cluster status is changed to indicate that autoscaling has stopped:

```
Status:
  Cluster State:
    Auto Scaling:
      Desired Replicas: 3
      Direction:        ScaleDown
      Metrics:
        Tps:
          Direction:      ScaleDown
          Time Threshold Crossed: 2024-02-02T21:24:17.690Z
        State:            Auto-scaling stopped by failure
```

At this point autoscaling has stopped. But more importantly, the elastic resizing *has not completed*. So it requires human intervention to either complete or revert the resizing before autoscaling can be resumed. Again, the first step is to determine what is causing the resize operation to fail and, if at all possible, correcting the situation. This is particularly true when resizing to reduce the size of the cluster because the only reliable way to recover the cluster is to complete the resizing operation.

To do this, you want to shutdown the cluster, then restart it specifying the number of replicas *before* the resize operation began. For example, if the cluster was resizing from 5 nodes to 3, you will want to restart the cluster with the number of replicas set to 5:

```
$ helm upgrade mydb voltdb/voltdb --reuse-values \
  --set cluster.clusterSpec.replicas=5
```

Finally, there are a few basic rules about what *not* to do with the cluster while using autoscaling:

- Do NOT manually initiate a resizing operation while autoscaling is enabled.
- Do NOT initiate a software upgrade while autoscaling is enabled.
- Do NOT shutdown the cluster while autoscaling is actively resizing the cluster.
- Do NOT disable autoscaling while autoscaling is actively resizing the cluster.

In all cases, it is safest to wait for all resizing to complete, then disable autoscaling by setting `cluster.clusterSpec.autoScaling.enabled` to *false* before performing any of the preceding actions.

3.4. Pausing and Resuming the Cluster

To pause the database — that is stop client activity through the client port when performing certain administrative functions — you set the property `cluster.clusterSpec.maintenanceMode` to true. For example, the following commands pause and then resume the database associated with release *mydb*:

```
$ helm upgrade mydb voltdb/voltdb \
  --reuse-values \
  --set cluster.clusterSpec.maintenanceMode=true
```

```
$ helm upgrade mydb voltdb/voltdb \
  --reuse-values \
  --set cluster.clusterSpec.maintenanceMode=false
```

3.5. Starting More than One Cluster Within a Namespace

By default, the Volt Helm charts assume there is only one cluster in each Kubernetes namespace. It is possible to run more than one Volt cluster within a namespace; however, to do so you need to start and stop the clusters and the operator separately. You do this by performing separate **helm install** operations for the operator and each cluster, using separate release names for each operation and setting the `cluster.enabled` and `operator.enabled` properties appropriately in each step. As with a regular helm install, you must specify the version of VoltDB you want to use when starting both the operator and the cluster.

For example, let's assume we want to start two clusters, *rome* and *venice*, in a single namespace. The steps for starting multiple Volt clusters in a single Kubernetes namespace are as follows:

1. Start the operator separately

Issue the helm install command setting the `operator.enabled` property to true and the `cluster.enabled` to false. Then wait for the operator to reach the ready state:

```
$ helm install voltoperator voltdb/voltdb \
  --set global.voltdbVersion=13.2.0 \
  --values opconfig.yaml \
  --set operator.enabled=true \
  --set cluster.enabled=false
```

Note that you can provide additional operator properties, separately or as a YAML file, as part of the install operation. See Section B.7, “Operator Configuration Options” for a list of available operator properties.

2. Start the first cluster

Once the operator is ready, you can start the first cluster, reversing the values for `operator.enabled` and `cluster.enabled` and providing whatever cluster-specific configuration you need:

```
$ helm install rome voltdb/voltdb \
  --set global.voltdbVersion=13.2.0 \
  --values romeconfig.yaml \
  --set cluster.clusterSpec.replicas=3 \
  --set operator.enabled=false \
  --set cluster.enabled=true
```

Again, wait for the pods of the cluster to reach the ready state before moving on to the next step.

3. Start subsequent clusters

Repeat step #2 for any other clusters you want to run in the namespace waiting after each install command for the pods to reach their ready state. In our example, we only have one other cluster:

```
$ helm install venice voltdb/voltdb \
  --set global.voltdbVersion=13.2.0 \
```

```
--values veniceconfig.yaml      \  
--set cluster.clusterSpec.replicas=3  \  
--set operator.enabled=false      \  
--set cluster.enabled=true
```

The key point when running multiple clusters within a single namespace is that there is only one Volt Operator and the operator executes one operation at a time. So be sure to wait for each Helm command to complete before issuing a new command. Because of the constraint to sequential processing in the Operator, we recommend limiting the number of simultaneous Volt clusters within any single namespace to three.

3.6. Stopping, Restarting, and Shutting Down Multiple Clusters Within a Namespace

Once you have multiple clusters running in the same namespace, you can stop and start the databases independently, the same way you would a single database, by setting the property `cluster.clusterSpec.replicas` to zero to stop the database and the correct number of nodes to restart it. For example, the following command stops the *rome* cluster without affecting the operator or other clusters in the namespace:

```
$ helm upgrade rome voltdb/voltdb      \  
--reuse-values                        \  
--set cluster.clusterSpec.replicas=0
```

If you want to shutdown and remove the clusters and operator entirely, you must first shutdown and delete the clusters, then delete the operator. The key point is that you cannot delete the Helm release for the operator until all of the releases it manages have been removed. Therefore, the process is:

4. Shutdown and delete the individual clusters

```
$ helm upgrade rome voltdb/voltdb --reuse-values  \  
--set cluster.clusterSpec.replicas=0  
$ helm upgrade venice voltdb/voltdb --reuse-values  \  
--set cluster.clusterSpec.replicas=0  
$ helm delete rome  
$ helm delete venice
```

5. Delete the operator

```
$      # Make sure all pods have been deleted  
$ kubectl get pods  
$      # Once all pods are gone, remove the Operator  
$ helm delete voltoperator
```

Chapter 4. Managing VoltDB Databases in Kubernetes

When running VoltDB in Kubernetes, you are implicitly managing two separate technologies: the database cluster — that consists of "nodes" and the server processes that run on them — and the collection of Kubernetes "pods" the database cluster runs on. There is a one-to-one relationship between VoltDB nodes and Kubernetes pods and it is important that these two technologies stay in sync.

The good news is that the VoltDB Operator and Helm manage the orchestration of Kubernetes and the VoltDB servers. If a database server goes down, Kubernetes recognizes that the corresponding pod is not "live" and spins up a replacement. On the other hand, if you *intentionally* stop the database without telling the Operator or Kubernetes, Kubernetes insists on trying to recreate it.

Therefore, whereas on traditional servers you use **voltadmin** and **sqlcmd** to manage both the cluster and the database content, it is important in a Kubernetes environment that you use the correct utilities for the separate functions:

- Use **kubectl** and **helm** to manage the cluster and the database configuration
- Use **voltadmin** and **sqlcmd** to manage the database contents.

The following sections explain how to access and use each of these utilities. Subsequent chapters explain how to perform common cluster and database management functions using these techniques.

4.1. Managing the Cluster Using kubectl and helm

The key advantage to using Kubernetes is that it automates common administrative tasks, such as making sure the cluster keeps running. This is because the VoltDB Operator and Helm charts manage the synchronization of VoltDB and Kubernetes for you. But it does mean you must use **helm** or **kubectl**, and *not* the equivalent **voltadmin** commands, to perform operations that affect Kubernetes, such as starting and stopping the database, resizing the cluster, changing the configuration, and so on.

When you start the database for the first time, you specify the VoltDB Helm chart and a set of properties that define how the cluster and database are configured. The result is a set of Kubernetes pods and VoltDB server processes known as a Helm "release".

To manage the cluster and database configuration you use the **helm upgrade** command to update the release and change the properties associated with the feature you want to control. For example, to change the frequency of periodic snapshots in the *mydb* release to 30 minutes, you specify the new value for the `cluster.config.deployment.snapshot.frequency` property, like so:

```
$ helm upgrade mydb voltdb/voltdb \
  --reuse-values \
  --set cluster.config.deployment.snapshot.frequency=30m
```

Note

It is also possible to use the **kubectl patch** command to change release properties, specifying the new property value and action to take as a JSON string. However, the examples in this book use the **helm upgrade** equivalent wherever possible as the helm command tends to be easier to read and remember.

One caveat to using the **helm upgrade** command is that it not only upgrades the release, it checks to see if there is a new version of the original chart (in this example, *voltdb/voltdb*) and upgrades that too. Problems could occur if there are changes to the original chart between when you first start the cluster and when you need to stop or resize it.

The public charts are not changed very frequently. But if your database is in production for an extended period of time it could be an issue. Fortunately, there is a solution. To avoid any unexpected changes, you can tell Helm to use a specific version of the chart — the version you started with.

First, use the **helm list** command to list all of the releases (that is, database instances) you have installed. In the listing it will include both the name and version of the chart in use. For example:

```
$ helm list
```

NAME	NAMESPACE	REVISION	UPDATED	STATUS	CHART	APP VERSION
mydb	default	1	2020-08-12 12:45:30	deployed	voltdb-1.0.0	10.0.0

You can then specify the specific chart version when you upgrade the release, thereby avoiding any unexpected side effects:

```
$ helm upgrade mydb voltdb/voltdb \
  --reuse-values \
  --set cluster.clusterSpec.replicas=7 \
  --version=1.0.0
```

4.2. Managing the Database Using **voltadmin** and **sqlcmd**

You manage the database using the VoltDB command line utilities **voltadmin** and **sqlcmd**, the same way you would in a traditional server environment. The one difference is that before you can issue VoltDB commands, you need to decide how to access the database cluster itself. There are two types of access available to you:

- Interactive access for issuing **sqlcmd** or **voltadmin** commands to manage the database
- Programmatic access, through the client or admin port, for invoking stored procedures

4.2.1. Accessing the Database Interactively

Kubernetes provides several ways to access the pods running your services. You can run commands on individual pods interactively through the **kubectl exec** command. You can use the same command to access the command shell for the pod by running **bash**. Or you can use port forwarding to open ports from the pods to your current environment.

In all three cases, you need to know the name of the pod you wish to access. When you start a VoltDB cluster with Helm, the pods are created with templated names based on the Helm release name and a sequential number. So if you named your three node cluster *mydb*, the pods would be called *mydb-voltdb-cluster-0*, *mydb-voltdb-cluster-1*, and *mydb-voltdb-cluster-2*. There are also separate pods for any auxiliary services, such as the Volt Management Center (VMC). If you are not sure of the names, you can use the **kubectl get pods** command to see a list:

```
$ kubectl get pods
```

NAME	READY	STATUS	RESTARTS	AGE
mydb-voltdb-cluster-0	1/1	Running	0	26m

mydb-voltdb-cluster-1	1/1	Running	0	26m
mydb-voltdb-operator-6bbb96b575-8z75x	1/1	Running	0	26m
mydb-voltdb-vmc-86c8d7b688-pnmlg	1/1	Running	0	26m

Having chosen a pod to use, running VoltDB commands interactively with **kubectl exec** is useful for issuing individual commands. After the command executes, **kubectl** returns you to your local shell. For example, you can check the status of the cluster using the **voltadmin status** command:

```
$ kubectl exec -it mydb-voltdb-cluster-0 -- voltadmin status
Cluster 0, version 10.0, hostcount 2, kfactor 0
 2 live host, 0 missing host, 0 live client, uptime 0 days 00:41:34.293
```

```
-----
HostId      Host Name
0mydb-voltdb-cluster-0
1mydb-voltdb-cluster-1
```

You can even use **kubectl exec** to start an interactive **sqlcmd** session, which stays active until you exit **sqlcmd**:

```
$ kubectl exec -it mydb-voltdb-cluster-0 -- sqlcmd
SQL Command :: localhost:21212
1> exit
$
```

Or you can pipe a file of SQL statements to **sqlcmd** as part of the command:

```
$ kubectl exec -it mydb-voltdb-cluster-0 -- sqlcmd < myschema.sql
```

However, **kubectl exec** commands execute in the context of the pod. So you cannot do things like load JAR files that are in your local directory. If you need to load schema and stored procedures, it is easier to use port forwarding, where ports on the pod are forwarded to the equivalent ports on localhost for your local machine, so you can run applications and utilities (such as **sqlcmd**, **voltdb**, and **voltadmin**) locally.

The **kubectl port-forward** command initiates port forwarding, which is active until you stop the command process. So you need a second process to utilize the linked ports. In the following example the user runs the voter sample application locally on a database in a Kubernetes cluster. To do this, one session enables port forwarding on the client port and the second session loads the stored procedures, schema, and then runs the client application:

Session #1

```
$ kubectl port-forward mydb-voltdb-cluster-0 21212
```

Session #2

```
$ cd ~/voltdb/doc/tutorials/voter
$ sqlcmd
SQL Command :: localhost:21212
1> load classes voter-procs.jar;
2> file ddl.sql;
3> exit
$ ./run.sh client
```

Port forwarding is useful for ad hoc activities such as loading schema and stored procedures to a running database and quick test runs of client applications. Port forwarding is *not* good for running production applications or any ongoing activities, due to its inherent lack of security or robustness as a network solution.

You can also use port forwarding to monitor the cluster using the web-based Volt Management Center (VMC) by forwarding port 8080 from the VMC service, using the service name. The following example also adds the `--address` argument so it is available to others on the local area network. (Otherwise it is only accessible as `localhost:8080` from the system on which the port forward command is issued.)

```
$ kubectl port-forward svc/mydb-voltdb-vmc 8080 --address=0.0.0.0
```

Note that there is only one instance of VMC for the entire cluster. By forwarding the port from the VMC service you can access all nodes of the cluster by using the **servers** menu on the DB Monitor tab.

4.2.2. Accessing the Database Programmatically

The approaches for connecting to the database interactively do not work for access by applications, because interactive access focuses on connecting to one node of the database. Applications are encouraged to create connections to *all* nodes of the database to distribute the workload and avoid bottle necks. In fact, the Java client for VoltDB has special settings to automatically connect to all available nodes (topology awareness) and direct partitioned procedures to the appropriate host (client affinity).

Kubernetes provides a number of services to make pods accessible beyond the Kubernetes cluster they run in; services such as cluster IPs, node ports, and load balancers. These services usually change the address and/or port number seen outside the cluster. And there are still other layers of networking and firewalls to traverse before these open ports are accessible outside of Kubernetes itself. This complexity, plus the fact that these services result in port numbers and external network addresses that do not match what the database itself thinks it is running on, make accessing the database from external applications impractical.

The recommended way to access a VoltDB database running in Kubernetes programmatically is to run your application as its own service within the same Kubernetes cluster as the database. This way you can take advantage of the existing VoltDB service names, such as `mydb-voltdb-cluster-client`, to connect to the database. You can then enable topology awareness in the Java client and let the client make the appropriate connections to the current VoltDB host IPs.

For example, if your database Helm release is called `mydb` and is running in the namespace `mydata`, the Java application code to initiate access to the database might look like the following:

```
org.voltdb.client.Client client = null;

ClientConfig config = new ClientConfig("", "");
config.setTopologyChangeAware(true);

client = ClientFactory.createClient(config);
client.createConnection("mydb-voltdb-cluster-client.mydata.svc.cluster.local");
```

Chapter 5. Updates and Upgrades

Once the database is up and running, Kubernetes works to keep it running in the configuration you specified. However, you may need to change that configuration as your database requirements evolve. Changes may be as simple as adding, deleting, or modifying database tables or procedures. Or you may want to modify the configuration of the database, adding new users, or even expanding the cluster by adding nodes.

The following sections describe some common update scenarios and how to perform them in a Kubernetes environment, including:

- Modifying the database schema
- Modifying the database or cluster configuration
- Upgrading the VoltDB software and Helm charts

5.1. Updating the Database Schema

Once the VoltDB database starts, you are ready to manage the database contents. Using Kubernetes does not change *how* you manage the database content. However, it does require a few extra steps to ensure you have access to the database, as described in Section 4.2.1, “Accessing the Database Interactively”.

First you need to identify the pods using the **kubect**l **get pods** command. You can then access the pods, individually, using the **kubect**l **exec** command, specifying the pod you want to access and the command you want to run. For example, to run `sqlcmd` on the first pod, use the following command:

```
$ kubect
```

l exec -it mydb-voltdb-cluster-0 -- sqlcmd
SQL Command :: localhost:21212
1>

You can execute a local batch file of `sqlcmd` commands remotely by piping the file into the utility. For example:

```
$ cat schema.sql
CREATE TABLE HELLOWORLD (
  HELLO VARCHAR(15), WORLD VARCHAR(15),
  DIALECT VARCHAR(15) NOT NULL
);
PARTITION TABLE HELLOWORLD ON COLUMN DIALECT;
$ kubect
```

l exec -it mydb-voltdb-cluster-0 -- sqlcmd < schema.sql
Command succeeded.
Command succeeded.
\$

Changing the database schema does not require synchronization with Helm or Kubernetes necessarily. However, if you specified the schema and/or procedure classes when you initially created the Helm release, it may be a good idea to keep those properties updated in case you need to re-initialize the database. (For example, when re-establishing a XDCR connection that was broken due to conflicts.) This can be done by updating the `cluster.config.schemas` and/or `cluster.config.classes` properties and their unique subproperties. For example:

```
$ helm upgrade mydb voltdb/voltdb \
  --reuse-values \
  --set-file cluster.config.schemas.mysql=schema.sql \
```

```
--set-file cluster.config.classes.myjar=procs.jar
```

Note that for the schema and classes you must specify a unique subproperty of your choosing for each file (In the previous example *mysql* and *myjar*). This way you can include multiple schema or class files by specifying each with a separate `--set-file` flag and a separate unique subproperty name (such as *sql1*, *sql2*, and so on).

5.2. Updating the Database Configuration

You can also change the configuration options for the database or the cluster while the database is running. In Kubernetes, you do this by updating the release properties rather than with the **voltadmin update** command.

How you update the configuration properties is the same for all p4roperties: you use the **helm upgrade** command to update the individual properties. However, what actions result from the update depend on the type of properties you want to modify:

- Dynamic database configuration properties that can be modified "on the fly" without restarting the database
- Static database configuration properties that require the database be restarted before they are applied
- Cluster configuration properties that alter the operation of the cluster and associated Kubernetes pods

The following sections describe these three circumstances in detail.

5.2.1. Changing Database Properties on the Running Database

There are a number of database configuration options that can be changed while the database is running. Those options include:

- Security settings, including user accounts

```
cluster.config.deployment.security.enabled  
cluster.config.deployment.users
```

- Import and export settings

```
cluster.config.deployment.export.configurations  
cluster.config.deployment.import.configurations
```

- Database replication connections settings

```
cluster.config.deployment.dr.connection
```

- Automated snapshots

```
cluster.config.deployment.snapshot.*
```

- Heartbeat timeout

```
cluster.config.deployment.heartbeat.timeout
```

- System settings:

```
cluster.config.deployment.systemsettings.compaction.*
cluster.config.deployment.systemsettings.elastic.*
cluster.config.deployment.systemsettings.flushinterval.*
cluster.config.deployment.systemsettings.query.timeout
cluster.config.deployment.systemsettings.resourcemonitor.*
```

For example, the following helm upgrade command changes the heartbeat timeout to 30 seconds:

```
$ helm upgrade mydb voltdb/voltdb \
  --reuse-values \
  --set cluster.config.deployment.heartbeat.timeout=30
```

When dynamic configuration properties are modified, the VoltDB Operator updates the running database configuration as soon as it is notified of the change.

5.2.2. Changing Database Properties That Require a Restart

Many database configuration properties are static — they cannot be changed without restarting the database. Normally, this requires manually performing a **voltadmin shutdown --save**, reinitializing and restarting the database cluster, then restoring the final snapshot. For example, command logging cannot be turned on or off while the database is running; similarly, the number of sites per host cannot be altered on the fly.

However, you *can* change these properties using the **helm upgrade** command and the VoltDB Operator will make the changes, but *not* while the database is running. Instead, the Operator recognizes the changes to the configuration, marks the database as requiring a restart, and then schedules a shutdown snapshot, reinitialization, and restart of the database for later.

For example, you cannot change the number of sites per host while the database is running. But the Operator does let you change the property in Kubernetes:

```
$ helm upgrade mydb voltdb/voltdb \
  --reuse-values \
  --set cluster.config.deployment.cluster.siteperhost=12
```

No action is taken immediately, since the change will require a restart and is likely to interrupt ongoing transactions. Instead, the Operator waits until you are ready to restart the cluster, which you signify by changing another property, `cluster.clusterSpec.allowRestartDuringUpdate`, to true:

```
$ helm upgrade mydb voltdb/voltdb \
  --reuse-values \
  --set cluster.clusterSpec.allowRestartDuringUpdate=true
```

Important

You should *never* set the `.allowRestartDuringUpdate` property to true for an XDCR cluster. If you do, the cluster will not be able to rejoin the XDCR mesh and will fail during startup. To change configuration options that require a restart for an XDCR cluster, you must:

- Permanently remove the cluster from the XDCR mesh and shut it down (Section 9.2, “Removing a Cluster Permanently”).
- Reinitialize and rejoin the cluster to the XDCR environment (Section 9.4, “Rejoining an XDCR Cluster That Was Previously Removed”).

If you are sure you are ready to restart the cluster when you change the configuration property, you can set the two properties at the same time so that the change takes immediate effect:

```
$ helm upgrade mydb voltdb/voltdb \
  --reuse-values \
  --set cluster.config.deployment.cluster.siteperhost=12 \
  --set cluster.clusterSpec.allowRestartDuringUpdate=true
```

Once `allowRestartDuringUpdate` is set to true, the Operator initiates the restart process, saving, shutting down, reinitializing, restarting and restoring the database automatically. Note that once the database is restarted, it is a good idea to reset `allowRestartDuringUpdate` to false to avoid future configuration changes triggering immediate restarts:

```
$ helm upgrade mydb voltdb/voltdb \
  --reuse-values \
  --set cluster.clusterSpec.allowRestartDuringUpdate=false
```

Warning

There are certain database configuration changes that cannot be made either on the fly or with a restart. In particular, do *not* attempt to change properties associated with directory paths or SSL configuration. Changing any of these properties will leave your database in an unstable state.

5.2.3. Changing Cluster Properties

There are properties associated with the environment that the VoltDB database runs on that you can also modify with the **helm upgrade** command. Most notably, you can increase the size of the cluster, using elastic scaling, by changing the `cluster.clusterSpec.replicas` property, as described in Section 3.3, “Resizing the Cluster with Elastic Scaling”.

Some properties affect the computing environment, such as environment variables and number of nodes. Others control the network ports assigned or features specific to Kubernetes, such as liveness and readiness. All these properties can be modified. However, they each have separate scopes that affect when the changes will go into effect.

Of particular note, pod-specific properties will not take effect until each pod restarts. If this is not a high availability cluster (that is, `K=0`), the Operator will wait until you change the property `cluster.clusterSpec.allowRestartDuringUpdate` to true before restarting the cluster and applying the changes. The same applies for any cluster-wide properties.

However, for a K-safe cluster, the Operator can apply pod-specific changes without any downtime by performing a *rolling upgrade*, stopping and replacing each pod in sequence. (You should only initiate such changes if the cluster is complete; that is, has no missing or stopped nodes.) For high availability clusters, the Operator will start applying pod-specific changes automatically via a rolling restart regardless of the `cluster.clusterSpec.allowRestartDuringUpdate` setting.

5.3. Upgrading the VoltDB Software and Helm Charts

When new versions of the VoltDB software are released they are accompanied by new versions of the Helm charts that support them. By default when you “install” a “release” of VoltDB with Helm, you get the latest version of the VoltDB software at that time. Your release will stay on its initial version of VoltDB as long as you don’t update the charts and VoltDB Operator in use.

You can upgrade an existing database instance to a recent version using a combination of **kubectl** and **helm** commands to update the charts, the operator, and the VoltDB software. The steps to upgrade the VoltDB software in Kubernetes are:

1. Update your copy of the VoltDB repository.
2. Update the custom resource definition (CRD) for the VoltDB Operator.
3. Upgrade the VoltDB Operator and software.

The following sections explain how to perform each step of this process, including a full example of the entire process in Example 5.1, “Process for Upgrading the VoltDB Software” However, when upgrading an XDCR cluster, there is an additional step required to ensure the cluster's schema is maintained during the upgrade process. Section 5.3.5, “Updating VoltDB for XDCR Clusters” explains the extra step necessary for XDCR clusters.

Note

To use the **helm upgrade** command to upgrade the VoltDB software, the starting version of VoltDB must be 10.1 or higher. See the *VoltDB Release Notes* for instructions when using Helm to upgrade earlier versions of VoltDB.

5.3.1. Updating Your Helm Repository

The first step when upgrading VoltDB is to make sure your local copy of the VoltDB Helm repository is up to date. You do this using the **helm repo update** command:

```
$ helm repo update
```

5.3.2. Updating the Custom Resource Definition (CRD)

The second step is to update the custom resource definition (CRD) for the VoltDB Operator. This allows the Operator to be upgraded to the latest version.

To update the CRD, you must first save a copy of the latest chart, then extract the CRD from the resulting tar file. The **helm pull** command saves the chart as a gzipped tar file and the **tar** command lets you extract the CRD. For example:

```
$ helm pull voltdb/voltdb
$ ls *.tgz
voltdb-3.1.0.tgz
$ tar --strip-components=2 -xzf voltdb-3.1.0.tgz \
    voltdb/crds/voltdb.com_voltdbclusters_crd.yaml
```

Note that the file name of the resulting tar file includes the chart version number. Once you have extracted the CRD as a YAML file, you can use it to replace the CRD in Kubernetes:

```
$ kubectl replace -f voltdb.com_voltdbclusters_crd.yaml
```

5.3.3. Upgrading the VoltDB Operator and Software

Once you update the CRD, you are ready to upgrade VoltDB. You do this using the **helm upgrade** command and specifying the new software version you wish to use on the command line. What happens when

you issue the **helm upgrade** command depends on whether you are performing a standard software upgrade or an in-service upgrade.

For a standard software upgrade, you simply issue the **helm upgrade** command specifying the software version in the `global.voltdbVersion` property. For example:

```
$ helm upgrade mydb voltdb/voltdb --reuse-values \
  --set global.voltdbVersion=13.2.1
```

When you issue the **helm upgrade** command, the operator saves a final snapshot, shuts down the cluster, restarts the cluster with the new version and restores the snapshot. For example, Example 5.1, “Process for Upgrading the VoltDB Software” summarizes all of the commands used to update a database release to VoltDB version 13.2.1.

Example 5.1. Process for Upgrading the VoltDB Software

```
$ # Update the local copy of the charts
$ helm repo update
$ # Extract and replace the CRD
$ helm pull voltdb/voltdb
$ ls *.tgz
voltdb-3.1.0.tgz
$ tar --strip-components=2 -xzf voltdb-3.1.0.tgz \
  voltdb/crds/voltdb.com_voltdbclusters_crd.yaml
$ kubectl replace -f voltdb.com_voltdbclusters_crd.yaml
$
$ # Upgrade the Operator and VoltDB software
$ helm upgrade mydb voltdb/voltdb --reuse-values \
  --set global.voltdbVersion=13.2.1
```

Note it is also possible to identify specific versions of the Volt Operator and/or chart when starting VoltDB. However, if you do, be sure to use matching versions of the Operator and chart to avoid any potential compatibility issues.

5.3.4. Using In-Service Upgrade to Update the VoltDB Software

Standard upgrades are convenient and can upgrade across multiple versions of the VoltDB software. However, they do require downtime while the cluster is shutdown and restarted. *In-Service Upgrades* avoid the need for downtime by upgrading the cluster nodes one at a time, while the database remains active and processing transactions.

To use in-service upgrades, you must have an appropriate software license (in-service upgrades are a separately licensed feature), the cluster must be K-safe (that is, have a K-safety factor of one or more), the cluster must be complete (that is, have no missing or stopped nodes), and the difference between the current software version and the version you are upgrading to must fall within the limits of in-service upgrades. The following sections describe:

- What versions can be upgraded using an in-service upgrade
- How to perform the in-service upgrade
- How to monitor the upgrade process
- How to rollback an in-service upgrade if the upgrade fails

5.3.4.1. The Scope of In-Service Upgrades

There are limits to which software versions can use in-service upgrades. The following table describes the rules for which releases can be upgraded with an in-service upgrade and which releases cannot.

✓ Patch Releases	You can upgrade between any two <i>patch releases</i> . That is, any two releases where only the third and final number of the version identifier changes. For example, upgrading from 13.1.1 to 13.1.4.
✓ Minor Releases	You can also use in-service upgrades to upgrade between two consecutive <i>minor releases</i>. That is where the second number in the version identifier differ. For example, you can upgrade from V13.2 to V13.3. You can also upgrade between any patch releases within those minor releases. For example, upgrading from V13.2.1 to V13.3.4. You <i>cannot</i> use in-service upgrades to upgrade more than one minor version at a time. In other words, you can upgrade from V13.2 to V13.3 but you cannot perform an in-service upgrade from V13.1 to V13.3. To transition across multiple minor releases your options are to perform consecutive in-service upgrades (for example, from V13.1.3 to V13.2.0, then from V13.2.0 to V13.3.8) or to perform a regular upgrade where all cluster nodes are upgrading at one time.
✓ Major Releases	Starting with VoltDB V13.3.1, you can even upgrade between two consecutive major releases using an in-service upgrade. For example, you can perform an in-service upgrade from V13 to V14. Upgrading across major versions requires that your starting version is an official Long-Term Support (LTS) release and the target version is either the initial release of the next major version or the LTS release of that version. For example, you can upgrade from the V13 LTS release 13.3.9 to either 14.0.x or 14.3.x (the LTS release) of version 14. You <i>cannot</i> use in-service upgrade to upgrade from a pre-LTS release (such as V13.1) or to an interim minor release (between n.0 and the n.LTS release) of the subsequent major version.

5.3.4.2. How to Perform an In-Service Upgrade

If your cluster meets the requirements, you can use the in-service upgrade process to automate the software update and eliminate the downtime associated with standard upgrades. The procedure for performing an in-service upgrade is:

1. Set the property `cluster.clusterSpec.enableInServiceUpgrade` to true to allow the upgrade.
2. Set the property `global.voltdbVersion` to the software version you want to upgrade to.

For example, the following command performs an in-service upgrade from V13.1.2 to V13.2.0:

```
helm upgrade mydb voltdb/voltdb --reuse-values \  
  --set cluster.clusterSpec.enableInServiceUpgrade=true \  
  --set global.voltdbVersion=13.2.0
```

5.3.4.3. Monitoring the In-Service Upgrade Process

Once you initiate an in-service upgrade, the process proceeds by itself until completion. At a high level you can monitor the current status of the upgrade using the `@SystemInformation` system procedure with the `OVERVIEW` selector and looking for the `VERSION` keyword. For example, in the following command

output, the first column is the host ID and the last column is the currently installed software version for that host. Once all hosts report using the upgraded software version, the upgrade is complete.

```
$ echo "exec @SystemInformation overview" | sqlcmd | grep VERSION
2 VERSION          13.1.2
1 VERSION          13.1.2
0 VERSION          13.1.3
```

During the upgrade, the Volt Operator reports various stages of the process as events to Kubernetes. So you can monitor the progression of the upgrade in more detail using the **kubectl get events** command. For example, the following is an abbreviated listing of events you might see during an in-service upgrade. (The messages often contain additional information concerning the pods or the software versions being upgraded from and to.)

```
$ kubectl get events -w
11m      Normal RollingUpgrade   mydb-voltdb-cluster   Gracefully terminating pod 2
11m      Normal RollingUpgrade   mydb-voltdb-cluster   Gracefully terminated pod 2
11m      Normal RollingUpgrade   mydb-voltdb-cluster   Recycling Gracefully terminated pod my
9m43s    Normal RollingUpgrade   mydb-voltdb-cluster   Recycled pod 2 has rejoined the cluste
9m42s    Normal RollingUpgrade   mydb-voltdb-cluster   Pod mydb-voltdb-cluster-2 is now READY
9m35s    Normal RollingUpgrade   mydb-voltdb-cluster   Gracefully terminating pod 1
[ . . . ]
```

Once the upgrade is finished, the Operator reports this as well:

```
5m10s    Normal RollingUpgrade   mydb-voltdb-cluster   RollingUpgrade Done.
```

5.3.4.4. Recovering if an Upgrade Fails

The in-service upgrade process is automatic on Kubernetes — once you initiate the upgrade, the Volt Operator handles all of the activities until the upgrade is complete. However, if the upgrade fails for any reason — for example, if a node fails to rejoin the cluster — you can *rollback* the upgrade, returning the cluster to its original software version.

The Volt Operator detects an error during the upgrade whenever the VoltDB server process fails. The failure is reported as an appropriate series of events to Kubernetes:

```
12m      Warning RollingUpgrade   mydb-voltdb-cluster   Rolling Upgrade failed upgrading from..
12m      Normal  RollingUpgrade   mydb-voltdb-cluster   Please update the clusterSpec image bac
```

In addition to monitoring the events, you may wish to use the **kubectl** commands **get events**, **get pods**, and **logs** to determine exactly why the node is failing. The next step is to cancel the upgrade by initiating a rollback. You do this by resetting the image tag to the original version number.

Invoking the rollback is a manual task. However, once the rollback is initiated, the Operator automates the process of returning the cluster to its original state. Consider the previous example where you are upgrading from V13.1.2 to V13.2.0. Let us assume three nodes had upgraded but a fourth was refusing to join the cluster. You could initiate a rollback by resetting the `global.voltdbVersion` property to V13.1.2:

```
$ helm upgrade mydb voltdb/voltdb --reuse-values \
--set global.voltdbVersion=13.1.2
```

Once you initiate the rollback, the Volt Operator stops the node currently being upgraded and restarts it using the original software version. After that process completes, the Operator goes through any node that had been upgraded, one at a time, downgrading them back to the original software. Once all nodes are reset and have rejoined the cluster, the rollback is complete.

Note that an in-service rollback can only occur if you initiate the rollback during the upgrade process. Once the in-service upgrade is complete and all nodes are running the new software version, resetting the image tag will force the cluster to perform a standard software downgrade, shutting down the cluster as a whole and restarting with the earlier version.

5.3.5. Updating VoltDB for XDCR Clusters

When upgrading an XDCR cluster, there is one extra step you must pay attention to. Normally, during the upgrade, VoltDB saves and restores a snapshot between versions and so all data and schema information is maintained. When upgrading an XDCR cluster, the data and schema is deleted, since the cluster will need to reload the data from another cluster in the XDCR relationship once the upgrade is complete.

Loading the data is automatic. But loading the schema depends on the schema being stored properly before the upgrade begins.

If the schema was loaded through the YAML properties `cluster.config.schemas` and `cluster.config.classes` originally and has not changed, the schema and classes will be restored automatically. However, if the schema was loaded manually or has been changed since it was originally loaded, you must make sure a current copy of the schema and classes is available after the upgrade. There are two ways to do this.

For both methods, the first step is to save a copy of the schema and the classes. You can do this using the **voltdb get schema** and **voltdb get classes** commands. For example, using Kubernetes port forwarding you can save a copy of the schema and class JAR file to your local working directory:

```
$ kubectl port-forward mydb-voltdb-cluster-0 21212 &
$ voltdb get schema -o myschema.sql
$ voltdb get classes -o myclasses.jar
```

Once you have copies of the current schema and class files, you can either set them as the default schema and classes for your database release before you upgrade the software or you can set them in the same command as you upgrade the software. For example, the following commands set the default schema and classes first, then upgrade the Operator and server software. Alternately, you could put the two `--set-file` and two `--set` arguments in a single command.

```
$ helm upgrade mydb voltdb/voltdb --reuse-values \
  --set-file cluster.config.schemas.mysql=myschema.sql \
  --set-file cluster.config.classes.myjar=myclasses.jar
$ helm upgrade mydb voltdb/voltdb --reuse-values \
  --set global.voltdbVersion=12.3.1
```

Chapter 6. Monitoring VoltDB Databases in Kubernetes

Once the database is running, you need to monitor the system to ensure reliable uptime and performance. Variations in usage, workload, or the operational environment can affect the dynamics of the data application, which may need corresponding adjustments to the schema, procedures, or hardware configuration. VoltDB provides system procedures (such as `@Statistics`) and the web-based Volt Management Center to help monitor current performance. But to provide persistent, historical intelligence concerning application performance it is best to use a dedicated metrics data store, such as Prometheus.

Prometheus is a metrics monitoring and alerting system that provides ongoing data collection and persistent storage for applications and other resources. By providing an open source industry standard for collecting and storing metrics, Prometheus allows you to:

- Offload monitoring from the database platform itself
- Combine metrics from VoltDB with other applications within your business ecosystem
- Query and visualize historical information about your database activity and performance (through tools such as Grafana)

Section 6.1, “Using Prometheus to Monitor VoltDB” explains how to configure your VoltDB database so the information you need is gathered and made available through Prometheus and compatible graphic consoles such as Grafana.

6.1. Using Prometheus to Monitor VoltDB

To monitor VoltDB with Prometheus on Kubernetes, you enable per pod metrics where each node of the cluster reports its own set of server-specific information. The servers make this data available in Prometheus format through an HTTP endpoint (`/metrics`) on the metrics port (which defaults to 11781). You can control the port number and other characteristics of the metrics system through Helm properties.

To enable Prometheus metrics, set the `cluster.config.deployment.metrics.enabled` property to `true`. You can also set the `cluster.serviceSpec.perpod.metrics.enabled` property to `true`, which creates a Kubernetes metrics service for each pod. Prometheus uses these metrics services to identify the Volt pods as targets for scraping. For example, the following command enables per pod metrics with default settings while initializing the `mydb` database cluster. It also sets the service type to `ClusterIP`:

```
$ helm install mydb voltdb/voltdb \
  --set-file cluster.config.licenseXMLFile=license.xml \
  --set cluster.clusterSpec.replicas=5 \
  --set cluster.config.deployment.metrics.enabled=true \
  --set cluster.serviceSpec.perpod.metrics.enabled=true \
  --set cluster.serviceSpec.service.metrics.type=ClusterIP
```

Once metrics are enabled, each Volt server reports its own information through the Prometheus endpoint on the metrics port. If you enable the per pod service, connection to the Prometheus server is handled automatically. If the service is not enabled or Prometheus is not configured to auto-detect targets, you will need to edit the Prometheus configuration to add the cluster nodes to the list of scraping targets.

Finally, if the database has security enabled, you will also need to configure Prometheus with the appropriate authentication information based on the truststore and password for the cluster. See the Prometheus documentation for more information.

Once Prometheus is scraping the Volt metrics, you can use tools such as Grafana to combine, analyze, and present the information in meaningful ways. There are example Grafana dashboards in the Volt Github repository (<https://github.com/VoltDB/volt-monitoring>) demonstrating some of the visualizations that are possible.

Chapter 7. Configuring Security in Kubernetes

There are two aspects to security with Volt Active Data — security within the database which is managed through user accounts and roles and network security between the database nodes, between the cluster and client applications, and between clusters in the case of cross datacenter replication (XDCR). For internal security, you define user accounts as part of the database configuration and assign them to roles that are defined as part of the schema. For network security, Volt recommends encryption and authentication certificates using the TLS/SSL protocol. The following sections explain how to configure both types of security within Kubernetes.

7.1. Configuring User Accounts and Roles Within The Database

User accounts allow you to control who has access to specific functions and procedures within the database. Security is enabled in the configuration with the `cluster.config.deployment.security.enabled` property. You must also use the properties to define the actual user names, passwords, and assigned roles. The `users` property expects a list of sub-elements so you must prefix each set of properties with a hyphen. The following example enables security and defines two user accounts:

```
cluster:
  config:
    deployment:
      security:
        enabled: true
      users:
        - name: controller
          password: superman
          roles: administrator
        - name: mitty
          password: thurber
          roles: user
```

7.1.1. Assigning Administrative Access to the Volt Operator

If you enable basic security, you must also tell the VoltDB operator which account to use when accessing the database. To do that, you define the `cluster.config.auth` properties, as shown below, which must specify an account with the built-in *administrator* role. The following example assigns the *controller* account for use by the VoltDB Operator:

```
cluster:
  config:
    auth:
      username: controller
      password: superman
```

7.1.2. Using Kubernetes Secrets to Store User Definitions

Normally, the definitions of user accounts and passwords are part of the overall database configuration options. To provide an additional level of security, you can store the user definitions separately in a Ku-

bernetes secret so they are not visible to operators or other personnel when starting or managing the database pods with Helm.

To hide the user definitions, you include them as a separate YAML file in a Kubernetes secret. You can name the secret whatever you like. But the secret must contain only one key and the key must be named `users.yaml`. The easiest way to do this is to create the secret using a file named `users.yaml`. If not, you must use the syntax `create secret --from-file=users.yaml={local-file-name}`.

The file in the secret must be a YAML file containing the user definitions and must start with the YAML file separator, three hyphens ("---"). For example:

```
$ cat users.yaml
---
cluster:
  config:
    deployment:
      users:
        - name: controller
          password: superman
          roles: administrator
        - name: mitty
          password: thurber
          roles: user
$ kubectl create secret generic secure-users-secret --from-file=users.yaml
```

Note that the `users.yaml` secret file contains only the user definitions. It does not contain the assignment of an administrators account to the Volt Operator. To hide the Operator assignment, you can create another Kubernetes secret with the keys `username` and `password`:

```
$ kubectl create secret generic operator-auth-secret \
  --from-literal=username=controller \
  --from-literal=password=superman
```

Once you create the user definition and Operator authorization secrets, you can start the cluster using those secrets by assigning them to the `cluster.config.usersSecretName` and `cluster.config.auth.credSecretName` properties, respectively. For example:

```
$ helm install mydb voltdb/voltdb \
  --values myconfig.yaml \
  --set cluster.config.usersSecretName=secure-users-secret \
  --set cluster.config.auth.credSecretName=operator-auth-secret \
  [ . . . ]
```

7.1.3. Updating User Account Secrets

If you define user accounts as part of the Helm configuration properties, you can update, add, or remove accounts like other Helm properties using the **helm upgrade** command. If you define the user accounts using a Kubernetes secret, you can update the account information either by updating the secret itself or by replacing the current secret with a new one with the updated information. Using the External Secrets Operator (ESO) simplifies the process of updating the secret in place. Or if you are managing your secrets manually, creating a new secret may be the easier approach. For example:

```
$ cat newusers.yaml
---
cluster:
```

```
config:
  deployment:
    users:
      - name: controller
        password: superman
        roles: administrator
      - name: jimmyolsen
        password: photography
        roles: user
$ kubectl create secret generic new-users-secret \
  --from-file=users.yaml=newusers.yaml
$ helm upgrade mydb voltdb/voltdb --reuse-values \
  --set cluster.config.usersSecretName=new-users-secret
```

Important

If you alter the account used to authorize access for the Volt Operator, you *must* update the `config.auth...` properties in the same **helm upgrade** command that updates the accounts.

7.2. Configuring TLS/SSL

Another important aspect of security is securing and authenticating the ports used to access the database. The most common way to do this is by enabling TLS/SSL to encrypt data and authenticate the servers using user-created certificates. The process for creating the private keystore and truststore using openssl is described in the section on "Creating the TLS/SSL Certificates" in the *Using VoltDB* guide. This process is the same whether you are running the cluster directly on bare metal servers or in Kubernetes.

Note

The following instructions describe creating and using TLS/SSL certificates in Privacy Enhanced Mail, or PEM, format. PEM is a text-based format that is easy to work with and is accepted by all Volt functions starting with Volt version 14.2.0 and Volt Operator 3.6.0. For earlier versions of Volt, certain functions require a JKS file instead of PEM. Please see the version 13 documentation for configuring TLS on Kubernetes for instructions on creating and using certificates acceptable to previous releases.

The following example uses YAML properties to enable TLS/SSL security, in much the same way you enable SSL encryption on bare metal. First you must enable SSL encryption, using the `cluster.config.deployment.ssl.enabled` property. Then you choose which ports will use SSL (in this example, the internal and external ports, but not XDCC). The YAML does *not* include the actual content of the truststore and keystore files. The example also chooses to have the operator authenticate the servers by setting `cluster.clusterSpec.ssl.insecure` to *false*.

```
cluster:
  config:
    deployment:
      ssl:
        enabled: true
        external: true
        internal: true
  clusterSpec:
    ssl:
      insecure: false
```

Note that the certificates themselves and associated key and trust stores are not specified in the YAML for the cluster deployment configuration. Instead, you pass the associated files to the Operator (and subsequently the server pods) in one of three ways:

- Using the `--set-file` argument to the Helm install
- Using Kubernetes secrets
- Using cert-manager

The following sections describe the three methods for configuring encryption. Two additional sections describe alternate methods for configuring TLS/SSL for the Operator and VMC and how to update the certificates when they expire.

7.2.1. Configuring TLS/SSL With the Helm Install `--set-file` Argument

Using the example YAML file in Section 7.2, “Configuring TLS/SSL” (and calling it `ssl.yaml`), we can complete the SSL configuration by specifying the truststore and keystore files on the **helm** command line with the `--set-file` argument:

```
helm install mydb voltdb/voltdb \
  --values myconfig.yaml \
  --values ssl.yaml \
  --set-file cluster.config.deployment.ssl.keystore.file=keystore.pem \
  --set-file cluster.config.deployment.ssl.truststore.file=truststore.pem
```

Two important notes concerning TLS/SSL configuration:

- If you enable SSL for the cluster's external interface and ports and you enable metrics, you must provide the appropriate SSL information to the Prometheus data retriever configuration so it can access the metrics port.
- If you enable SSL for the cluster, you must repeat the specification of the truststore and keystore files every time you update the configuration. Using the `--reuse-values` argument on the **helm upgrade** command is *not* sufficient.

7.2.2. Using Kubernetes Secrets to Store and Reuse TLS/SSL Information

An alternative method is to store the key and trust stores in a Kubernetes secret. Secrets are a standard feature of Kubernetes that allow you to store sensitive information as key value pairs in a protected space. Three advantages of using a secret are:

- You do not have to enter sensitive TLS/SSL information in plain text when configuring or updating your database.
- The secret is used automatically for subsequent updates; you do not have to repeatedly specify the TLS/SSL files when updating the database configuration.
- You can reuse the same secret for multiple database instances and services.

To use a Kubernetes secret to store the TLS/SSL information for your database, you must first create the necessary PEM files as described in the section on “Creating the TLS/SSL Certificates” in the *Using Volt-*

DB guide. Next you create your Kubernetes secret using the **kubect** **create secret** command, specifying the key names and corresponding artifacts as arguments. For example:

```
$ kubectl create secret generic my-ssl-creds \
  --from-file=keystore_data=keystore.pem \
  --from-file=truststore_data=truststore.pem
```

It is critical you use the key names *keystore_data* and *truststore_data* (as well as *keystore_password* and *truststore_password* if you created a password for the certificate). If not, the Volt Operator will not be able to find them. Also, the secret must be in the same Kubernetes namespace as the Helm release you are configuring.

Once you create the secret you can use it to configure your database by *not* setting any of standard SSL properties such as the `cluster.config.deployment.ssl...` properties or `cluster.clusterSpec.ssl.certificateFile`. Instead, set the property `cluster.config.deployment.t.ssl.sslSecret.certSecretName`. Using the secret created in the preceding example, the configuration of your database will look something like this:

```
cluster:
  config:
    deployment:
      ssl:
        sslSecret:
          certSecretName: my-ssl-creds
```

7.2.3. Using Kubernetes cert-manager to Store TLS/SSL Certificates

Another alternative for maintaining the TLS/SSL information is to use the Kubernetes cert-manager (cert-manager.io). The *cert-manager* is an add-on for Kubernetes that helps you create and maintain certificates and other private information in Kubernetes. If you wish to use cert-manager for self-signed certificates, you not only use it to store the certificate and truststore, you create them with cert-manager as well. (For more detailed information concerning cert-manager, see the cert-manager documentation.)

The basic steps for storing self-signed TLS/SSL credentials in cert-manager are:

1. Create an *issuer* resource in Kubernetes that will generate and authenticate the certificate. You only need to do this once for the namespace and multiple certificate requests can use the same issuer.
2. Create a *request* for the issuer to generate the actual TLS/SSL certificate and store it in a Kubernetes secret.
3. Specify the resulting certificate secret in the VoltDB configuration and start your cluster.

You create the cert-manager issuer and the certificate request using YAML properties. The easiest way to do this is by typing the property declarations into a YAML file. For example, the following two YAML files create a cert-manager issuer service and request a certificate.

create-issuer.yaml

```
apiVersion: cert-manager.io/v1
kind: Issuer
metadata:
  name: selfsigned-issuer
  namespace: mydb
```

```
spec:
  selfSigned: {}

request-cert.yaml

apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
  name: my-ssl-certificate
  namespace: mydb
spec:
  commonName: voltdb.com
  duration: 8766h
  secretName: my-ssl-creds
  issuerRef:
    name: selfsigned-issuer
    kind: Issuer
  privateKey:
    algorithm: RSA
    encoding: PKCS8
    size: 2048
  usages:
    - server auth
```

Three key points to note about the certificate request are:

- The issuer must be in the same namespace as the database that uses the certificate.
- You specify the duration of the certificate in hours. In this example, 8766 hours, or one year.
- The encoding must be PKCS8

Once you create the YAML files, you can create the issuer and request the certificate:

```
$ kubectl apply -f create-issuer.yaml    # Do only once
$ kubectl apply -f request-cert.yaml
```

Finally, in your database configuration, you point to the secret created by the certificate request (in this case, *my-ssl-creds*) the same way you would for a manually created secret:

```
cluster:
  config:
    deployment:
      ssl:
        sslSecret:
          certSecretName: my-ssl-creds
```

7.2.4. Using Certificate Revocation Lists (CRLs)

A certificate revocation list (CRL) lets you deny access to the server from certain clients based on their certificates, when using mutual TLS (mTLS). For example, say you had separate client certificates for each vendor with access to your database. If the partnership with a vendor ends, you want to now deny access to the vendor's clients by adding their certificate to the list of CRLs. Note that the CRL files must be in PEM format. See the section on Using CRLs in the *Using VoltDB* guide for instructions on converting files in other formats to PEM using the **openssl** utility.

In Kubernetes, you add CRLs to the server in one of two ways:

- By specifying the CRLs, one at a time, using the `--set-file` qualifier when installing the database release. For example:

```
helm install mydb voltdb/voltdb
  --values myconfig.yaml
  --set-file cluster.config.deployment.ssl.crl.file.a=/etc/ssl/local/crl/crl1.pem
  --set-file cluster.config.deployment.ssl.crl.file.b=/etc/ssl/local/crl/crl2.pem
  --set-file cluster.config.deployment.ssl.crl.file.c=/etc/ssl/local/crl/crl3.pem
```

The last name of the `cluster.config.deployment.ssl.crl.file.{name}` property can be any name you choose as long as it is unique within the set of CRL file property names.

- By creating a configuration map containing the CRLs. Create a directory with one or more CRL files in it. Then create a configuration map using the **kubectl create configmap** command pointing to the CRL directory. For example, using the directory name `/etc/ssl/local/crl` from the preceding example and naming the configuration map `mycrlmap`:

```
kubectl create configmap mycrlmap --from-file /etc/ssl/local/crl
```

Once you create the configuration map, you can add it to the Helm configuration using the `cluster.config.deployment.ssl.crl.configName` property when installing the database release:

```
helm install mydb voltdb/voltdb
  --values myconfig.yaml
  --set cluster.config.deployment.ssl.crl.configName=mycrlmap
```

7.2.5. Configuring TLS/SSL for the Volt Operator and VMC

By default, when you enable TLS/SSL, the Volt Operator and the Volt Management Console (VMC) automatically use the truststore configured for the servers as their truststore for authenticating the server. In other words, you do not have to set any SSL configuration properties yourself for these two services to the servers. Of course, you can choose to explicitly identify a separate truststore for these other services and provide the associated files yourself, if you wish. But by default, configuring SSL for the server automatically configures the Operator and VMC correctly to respond to the encrypted messages.

The same is true if you enable mutual TLS (mTLS) for the cluster. In this case, the Operator and VMC use the same certificate (and truststore) as the servers by default. If you want to have these services use a distinct client certificate, you need to specify that certificate using the appropriate YAML properties for those services. For example, the following Helm command loads a separate certificate for the Volt Operator and uses the root truststore for both the server and Operator:

```
helm install mydb voltdb/voltdb
  --values myconfig.yaml
  --values ssl.yaml
  --set-file cluster.config.deployment.ssl.keystore.file=server.keystore.pem
  --set-file cluster.config.deployment.ssl.truststore.file=root.truststore.pem
  --set-file cluster.clusterSpec.ssl.clientCertFile=client.keystore.pem
```

However, the Jetty service that VMC uses to provide a public interface to web browsers and for the JSON programming interface does *not* support PEM-formatted files. Therefore to enable TLS/SSL between the VMC service and web browsers and client applications using the JSON API, you should use either JKS or PKCS12 formatted key and trust stores. You can either convert an existing PEM file to JKS or PKCS12

format using the **openssl** utility or see the VoltDB V13 documentation for instructions on creating user-defined certificates using the **keytool** utility. (To create PKCS12 rather than JKS files, simply replace the `-storetype jks` qualifier with `-storetype pkcs12` in the examples.) Once you create the necessary keystore and truststore files, use the instructions in Section 7.2.2, “Using Kubernetes Secrets to Store and Reuse TLS/SSL Information” to create a Kubernetes secret to store the files and then assign it to the `vmc.service.ssl.sslSecret.certSecretName` property.

7.3. Updating TLS/SSL Security Certificates

TLS certificates have an expiration date. If you are using TLS/SSL to encrypt data (either internally, externally, or both), you will need to update those certificates *before* they expire to ensure minimal disruption to normal operation.

One of the advantages of using cert-manager to create and manage your certificates is that it automatically updates the certificates before they expire. If you are not using cert-manager — that is, you are either creating your own secret to contain the keystore and truststore or defining them manually with helm properties — you will need to update the certificates yourself. Either way, shortly after the certificates are updated in Kubernetes, the operator takes responsibility for applying the new credentials to the cluster, the Operator, and the auxiliary services as appropriate.

To update the TLS keystores, truststores, and credentials when using a self-defined secret, you must:

1. Create a new version of the truststore and keystore using a certificate with a new expiration date.
2. Delete the current Kubernetes secret.
3. Create a new version of the same secret using the new files.

You create the new truststore and keystore using the same **openssl** commands used to create the original files, as described in Section 7.2, “Configuring TLS/SSL”. You then update the secret by deleting and recreating the secret using the **kubectl create secret** command from earlier, making sure you use the same name for the secret but the new SSL files. For example:

```
$ kubectl delete secret/my-ssl-creds
$ kubectl create secret generic my-ssl-creds \
  --from-file=keystore_data=newkey.pem \
  --from-file=truststore_data=newtrust.pem
```

If you defined the TLS/SSL credentials manually using Helm properties, you will need to reapply the new truststore and keystore files using a **helm upgrade** command and the `--set-file` flag.

Chapter 8. Cross Datacenter Replication in Kubernetes

Previous chapters describe how to run a single VoltDB cluster within Kubernetes. Of course, you can run multiple independent VoltDB databases in Kubernetes. You do this by starting each cluster in separate regions, under different namespaces within the same Kubernetes cluster, or running a single instance of the VoltDB Operator managing multiple clusters in the same namespace. However, some business applications require the same database running in multiple locations — whether for data redundancy, disaster recovery, or geographic distribution. In VoltDB this is done through *Cross Datacenter Replication*, or XDCR.

Important

Please note that in addition to the guidance specific to Kubernetes provided in this chapter, the following rules apply to XDCR in *any* operating environment:

- **You must have command logging enabled** for three or more clusters.
- **You can only join (or rejoin) one cluster at a time** to the XDCR environment.

Command logging is always recommended when using XDCR to ensure durability. Using XDCR without command logging on two clusters, it is possible for transactions processed on one cluster to be lost if the cluster crashes before the binary log is sent to the other cluster. However, for three or more clusters, command logging is *required*. Without command logging, not only can XDCR transactions be lost, but *it is likely the databases will diverge without warning*, if a cluster crashes after sending a binary log to one cooperating cluster but not to the other.

8.1. Requirements for XDCR in Kubernetes

Once established, XDCR in Kubernetes works the same way it does in any other network environment, as described in the chapter on Database Replication in the *Using VoltDB* guide. The key difference when using XDCR in Kubernetes is how you establish the initial connection between the clusters. Unlike traditional servers with known IP addresses, in Kubernetes network addresses are assigned on the fly and are not normally accessible outside individual namespaces or regions. Therefore, you must do additional work to create the appropriate network relationships. Specifically, you must:

- **Establish a network mesh** between the Kubernetes clusters containing the VoltDB databases so that the nodes of each VoltDB cluster can identify and resolve the IP addresses and ports of all the nodes from the other VoltDB clusters.
- **Configure the VoltDB clusters**, including properties that identify the type of mesh involved and mesh-specific annotations that determine what network addresses and ports to use.

The following sections describe the different approaches to establishing a network mesh and how to configure the clusters in each case.

8.2. Choosing How to Establish a Network Mesh

For XDCR to work, each cluster must be able to identify and connect to the nodes of the other cluster. Establishing the XDCR relationship occurs in two distinct phases:

1. **Network Discovery** — First, the clusters connect over the replication port (port 5555, by default). The initial connection confirms that the configurations are compatible, that the schema of the two clusters match for all DR tables, and that there is data in only one of the clusters.
2. **Replication** — Once the clusters agree on the schema, each cluster sends a list of node IP addresses and ports to the other cluster and multiple connections are made, node-to-node, between the two clusters. If there is existing data, a synchronization snapshot is sent between the clusters and then replication begins.

For the network discovery phase, each cluster must have a clearly identifiable network address that the other cluster can specify as part of its XDCR configuration. For the replication phase, each cluster must have externally reachable network addresses for each node in the cluster that it can advertise during the discovery phase and that the other cluster uses to make the necessary connections for replication.

Since, by default, the ports on a Kubernetes pod are not externally accessible, you must use additional services to make the VoltDB nodes accessible. Three such options are:

- **Kubernetes Load Balancers** — One way to establish a network mesh is to use the built-in load balancer service within Kubernetes. Load balancers provide a defined, persistent external interface for internal pods. The advantage of using load balancers is that they are a native component of Kubernetes and are easy to configure. The disadvantage is that if you are running your VoltDB clusters in a hosted environment, load balancers tend to be far more expensive than regular pods and creating a separate load balancer for each node in the cluster to handle the replication phase can be prohibitively expensive unless you are managing your own infrastructure.
- **Kubernetes Node Ports** — An alternative to load balancers is using node ports. Node ports, like load balancers, are native services of Kubernetes and provide an externally accessible interface for the internal pods. However, unlike load balancers where the addresses are persistent over time, node ports take on the addresses of the underlying Kubernetes nodes and therefore can change as Kubernetes nodes are recycled. Therefore node ports are not appropriate for the Network Discovery phase. On the other hand, they can be a cheaper alternative to load balancers for the replication phase, since the cluster can advertise the current set of node port addresses as pods come and go.
- **Network Mesh Services** — These additional services, such as Consul, create a network mesh between Kubernetes clusters and regions. They essentially act as a virtual private network (VPN) within Kubernetes so the VoltDB clusters can interoperate as if they were local to each other. The advantage of using network mesh services is that configuring the VoltDB clusters is simpler, since all of the network topology is handled separately. The deficit is that this requires yet another service to set up. And the configuration of these services can be quite complex, requiring a deep understanding of — and access to — the networking layer in Kubernetes.

Which networking solution you use is up to you. You can even mix and match the alternatives — using, for example, a single load balancer per cluster for the Network Discovery phase and individual node ports for each VoltDB cluster node during the replication phase.

Another alternative, if your pods have a secondary network interface (through a service such as Multus), is having each cluster advertise its replication address directly from that interface, rather than from a per-pod load balancer. You do this with the `cluster.clusterSpec.network.dr` property, described in Section 2.3.2, “Configuring Multiple Network Interfaces”. This is an alternative to the per-pod load balancer and `publicIPFromService` approach.

Note

The `network.dr` property takes precedence over `publicIPFromService`. When `network.dr` is configured, the DR address is resolved on the pod, so the Operator does not wait for

the DR load balancer to be assigned an external IP — even if `publicIPFromService` is also set. If you set both, the Operator logs a warning and ignores `publicIPFromService` for DR.

You define the type of network mesh to use and how to connect using YAML properties when you configure your clusters. In general, the Helm properties starting with `cluster.config.deployment.dr`, such as `id` and `role`, are generic properties common to all XDCR implementations. Helm properties starting with `cluster.serviceSpec` define the type of network mesh to use and annotations specific to the network type.

The following sections explain how to configure XDCR using Helm properties, with individual sections discussing the differences necessary for various networking options, including:

- Common XDCR Properties
- Configuring XDCR in Local Namespaces
- Configuring XDCR Using Load Balancers
- Configuring XDCR Using Node Ports for Replication
- Configuring XDCR Using Network Services

8.3. Common XDCR Properties

No matter what approach you choose for establishing the network mesh, you must first configure the clusters as members of the XDCR quorum the same way you do on bare metal. That is, you must assign:

- A unique DR ID for each cluster between 0 and 127
- The cluster role (XDCR)
- At least one node from the other cluster as the point of connection for the Network Discovery phase

The following table configures XDCR using DR ID 1 with a connection to the cluster with a release name of *brooklyn*.

```
cluster:
  config:
    deployment:
      dr:
        id: 1
        role: xdcr
        connection:
          enabled: true
          source: \
"brooklyn-voltdb-cluster-dr:5555"
```

8.4. Configuring XDCR in Local Namespaces

The easiest way to configure XDCR clusters is when the VoltDB clusters are within the same Kubernetes namespace or cluster. In this case, the cluster IP addresses are all locally visible and so do not need any additional network setup. The first step is to enable the DR service using the `cluster.serviceSpec.dr.enabled` property:

```
cluster:
  serviceSpec:
```

```
dr:
  enabled: true
```

Next, you must provide the address of a replication port from one node of the remote cluster as the `source` property.

In Kubernetes the cluster nodes are assigned unique host names based on the initial Helm release name (that is, the name you assigned the cluster when you installed it). The VoltDB Operator also creates services that abstract the individual server addresses and provide a single entry point for specific ports on the database cluster. The two services of interest are DR and client, which will direct traffic to the corresponding port (5555 or 21212 by default) on an arbitrary node of the cluster. If the two database instances are within the same Kubernetes cluster, you can use the DR service to make the initial connection between the database systems, as shown in the following YAML configuration file.

If the databases are running in different namespaces, you will need to specify the fully qualified service name as the connection source in the configuration, which includes the namespace. So, for example, if the *manhattan* database is in namespace *ny1* and *brooklyn* is in *ny2*, the YAML configuration files related to XDCR for the two clusters would be the following.

Manhattan Cluster

```
cluster:
  config:
    deployment:
      dr:
        id: 1
        role: xdcr
        connection:
          enabled: true
          source: "brooklyn-voltdb-cluster-dr.ny2.svc.cluster.local:5555"
```

Brooklyn Cluster

```
cluster:
  config:
    deployment:
      dr:
        id: 2
        role: xdcr
        connection:
          enabled: true
          source: "manhattan-voltdb-cluster-dr.ny1.svc.cluster.local:5555"
```

8.5. Configuring XDCR Using Load Balancers

Kubernetes load balancers are an alternative for making VoltDB clusters accessible outside the Kubernetes cluster or region they are in. In this case you are not using load balancers for their traditional role, balancing the load between multiple pods. Instead, the load balancers are solely used to provide externally accessible IP addresses.

There are two approaches to using load balancers. The first approach is to assign a load balancer for each node of the cluster. Since the nodes are externally reachable through persistent IP addresses on their corresponding load balancer, the load balancers can be used for both the network discovery and replication phases. The second approach is to use only one load balancer for the entire cluster to provide network discovery, and use virtual network peering, available from your hosting provider, for replication.

Many hosting platforms, such as Google Cloud or AWS, provide proprietary mechanisms for performing network peering between regions or data centers. Each of these solutions has its own unique set up and configuration, separate from the configuration of VoltDB and the VoltDB Operator. As a result, using a network peering service is not as simple as the use of load balancers for replication. However, they can be significantly more cost effective when paired with a single load balancer for network discovery.

There is also the choice of assigning the IP addresses for the load balancers dynamically, or having them selected from a range of static addresses. Dynamic assignment is simpler, since you do not need to arrange with your hosting provider for pre-assigned IPs or hostnames. However, dynamic addresses also mean you do not know what the addresses are *until the cluster starts*. This means the remote XDCR cluster cannot assign the `source` property until after the cluster starts with its associated load balancers and you can determine the IP addresses assigned to them.

8.5.1. Separate Load Balancers For Each Node (`cluster.serviceSpec.perpod`)

First you must assign the `DR id` and `role` as Helm properties. If the remote cluster is using static addresses, you can specify one of its nodes as the `source`, as in the following example. If you are using dynamic load balancers, leave the `source` property blank and use the **`helm upgrade --set`** command once the clusters are running to assign a resulting node address for the remote cluster.

```
cluster:
  config:
    deployment:
      dr:
        id: 1
        role: xdcr
        connection:
          enabled: true
          source: "chicago-dc-2" # Remote cluster
```

Then in the `cluster.serviceSpec` section, you enable `perpod` by setting its `type` to `LoadBalancer`. You will also want to set the `dr.enabled` property to `true` so the per pod load balancers are used for network discovery as well as replication.

For dynamically assigned addresses, set the `publicIPFromService` to `true`:

```
cluster:
  serviceSpec:
    perpod:
      type: LoadBalancer
      publicIPFromService: true
    dr:
      enabled: true
```

For static IP addresses, use the `staticIPs` property to specify the addresses to assign when creating the load balancers and, again, set `dr.enabled` to `true`.

```
cluster:
  serviceSpec:
    perpod:
      type: LoadBalancer
      staticIPs:
        - 12.34.56.78
```

```
- 12.34.56.79
- 12.34.56.80
dr:
  enabled: true
```

8.5.2. Single Load Balancer For Discovery with Virtual Networking Peering (cluster.serviceSpec.dr)

To reduce the number of resources needed to connect XDCR clusters in different regions, you can use a single load balancer for network discovery and use virtual network peering services from your hosting provider for connecting the two clusters during replication. How you set up and configure your network peering is specific to each provider. See your provider's documentation for additional information. This section describes how to set up a single Kubernetes load balancer for network discovery once you have your network peering established.

First you must assign the `DR id` and `role` as Helm properties and, if known in advance, the `source` for the remote cluster:

```
cluster:
  config:
    deployment:
      dr:
        id: 1
        role: xdcr
        connection:
          enabled: true
          source: "chicago-dc-2" # Remote cluster
```

Then in the `cluster.serviceSpec` section, you enable the `dr` service (rather than `perpod`) and set its `type` to `LoadBalancer`. You may also need to provide additional annotations that help configure the service. These annotations are specific to the host environment you are using. So, for example, the following configuration provides annotations for AWS and the Google Cloud:

```
cluster:
  serviceSpec:
    dr:
      enabled: true
      type: LoadBalancer
      annotations:
        # Google Cloud
        networking.gke.io/load-balancer-type: "Internal"
        networking.gke.io/internal-load-balancer-allow-global-access: "true"

        # AWS
        service.beta.kubernetes.io/aws-load-balancer-internal: "true"
        service.beta.kubernetes.io/aws-load-balancer-type: "nlb"
```

8.6. Configuring XDCR Using Node Ports for Replication

Kubernetes node ports are another option for providing external access to the VoltDB cluster for replication. Node ports are similar to load balancers in that they provide an externally accessible network address

for individual ports. Node ports are different in that the addresses are transitory — the address and/or port number will change as pods come and go. So node ports are less practical for the Network Discovery phase. However, they can be a cheap alternative for providing external access during the replication phase, since the cluster can advertise the new addresses as its topology changes.

It is also possible to mix and match solutions. So a single load balancer can be used to provide the Network Discovery service for a cluster, while node ports provide per pod network addresses for the replication phase, as described next.

Again, you start by assigning the `DR id` and `role` as Helm properties and, if known in advance, the `source` for the remote cluster:

```
cluster:
  config:
    deployment:
      dr:
        id: 1
        role: xdcr
        connection:
          enabled: true
          source: "chicago-dc-2" # Remote cluster
```

You then define the load balancer for Network Discovery by setting the values of the `cluster.serviceSpec.dr` properties `enabled` to `true` and `type` to `LoadBalancer`.

```
cluster:
  serviceSpec:
    dr:
      enabled: true
      type: LoadBalancer
```

Finally, define the replication phase as using node ports by configuring `cluster.serviceSpec.perpod` properties `type` to `NodePort` and `dr.enabled` to `true`. You can also use the `dr.startReplicationNodePort` property to specify the starting port number for the externally accessible ports assigned to the node ports.

```
cluster:
  serviceSpec:
    perpod:
      type: NodePort
    dr:
      enabled: true
      startReplicationNodePort: 33111
```

8.7. Configuring XDCR Using Network Services

The goal of network services, such as Consul, is to make Kubernetes pods in different clusters or regions appear as if they were local to each other. This makes configuring XDCR within VoltDB itself easier; in most cases it is almost identical to how you configure clusters within local namespaces. However, how you configure the network service itself is very dependent on which service you are using and the hosting environment in which you are operating.

Using Consul as an example, Consul provides a "sidecar" — an additional process running in the same pod as the VoltDB process — that makes remote pods and clusters appear to be local to the pod itself.

So rather than providing a remote IP address and port as the source for XDCR Network Discovery, you specify a local port. For example:

```
cluster:
  config:
    deployment:
      dr:
        id: 1
        role: xdcr
        connection:
          enabled: true
          source: "localhost:4444"
```

What port you specify and how you configure and start Consul and the Consul sidecar, is specific to the Consul product and your implementation of it. The same is true when using other third-party networking services. You may also need to provide additional annotations within the Helm configuration to complete the network setup, depending upon which network service you use. For example:

```
cluster:
  clusterSpec:
    additionalAnnotations:
      "consul.hashicorp.com/connect-service": "chicago-voltdb-cluster"
      "consul.hashicorp.com/connect-service-upstreams": "chicago-voltdb-cluster:55"
```

See the product documentation for the specific service for further information.

Chapter 9. Managing XDCR Clusters in Kubernetes

Once you have configured your XDCR clusters and your network environment, you are ready to start the clusters. You begin by starting two of the clusters. (Remember, only one of the clusters can have data in the DR tables before the XDCR communication begins.) Once the schema of the DR tables in two databases match, synchronization starts. After the initial two databases are synchronized, you can start additional XDCR clusters, one at a time.

There are several management procedures that help keep the clusters in sync, especially when shutting down or removing clusters from the XDCR environment. In other environments, these procedures use **voltadmin** commands, such as **shutdown** and **dr reset**. In Kubernetes, you execute these procedures through the VoltDB Operator using Helm properties. Activities include:

- Removing a cluster temporarily
- Removing a cluster permanently
- Resetting XDCR when a cluster is lost
- Rejoining a cluster that was removed

9.1. Removing a Cluster Temporarily

If you want to remove a cluster from the XDCR environment temporarily, you simply shutdown the cluster normally, by setting the number of replicas to zero. This way, when the cluster restarts, the command logs will take care of recovering all of the data and re-establishing the XDCR "conversations" with the other clusters:

```
--set cluster.clusterSpec.replicas=0
```

9.2. Removing a Cluster Permanently

If you want to remove a cluster from the XDCR environment permanently, you want to make sure it sends all of its completed transactions to the other clusters before it shuts down. You do this by setting the DR role to "none" to perform an orderly shutdown:

```
--set cluster.config.deployment.dr.role="none"  
--set cluster.clusterSpec.replicas=0
```

Of course, you do not have to shut the cluster down immediately. You can simply remove it from the XDCR environment. Note that if you do so, the data in the current cluster will diverge from those clusters still participating in XDCR. So only do this if you are sure you want to maintain a detached copy of the data:

```
--set cluster.config.deployment.dr.role="none"
```

Setting the role to "none" drops the cluster out of the XDCR mesh. However, it does not actually change the configuration property on the cluster itself until the cluster restarts. In other words, there is a pending configuration change. So if the property `cluster.clusterSpec.allowRestartDuringUpdate` is set to "TRUE", not only will the cluster drop out of the XDCR mesh, it will automatically restart.

Finally, if you cannot perform an orderly removal from XDCR — for example, if one of the other clusters is offline or if sending the outstanding transactions will take too long and you are willing to lose that data

— you can set the property `cluster.clusterSpec.dr.forceDrop` to "TRUE" to force the cluster to drop out of the XDCR mesh without finalizing its XDCR transfers. Once the cluster has been removed, it is advisable to reset this property to "FALSE" so future procedures revert to the orderly approach of flushing the queues.

```
--set cluster.clusterSpec.dr.forceDrop=TRUE
--set cluster.config.deployment.dr.role="none"
--set cluster.clusterSpec.replicas=0
...
--set cluster.clusterSpec.dr.forceDrop=FALSE
```

9.3. Resetting XDCR When a Cluster Leaves Unexpectedly

Normally, when a cluster is removed from XDCR in an orderly fashion, the other clusters are notified that the cluster has left the mesh. However, if a cluster leaves unexpectedly — for example, if it crashes or is shutdown and deleted without setting its role to "none" to notify the other clusters — the XDCR network still thinks the cluster is a member and may return. As a result, the remaining clusters continue to save DR logs for the missing member, using up unnecessary processing cycles and disk space. You need to reset the XDCR network mesh to correct this situation.

To reset the mesh you notify the remaining clusters that the missing cluster is no longer a member. You do this by adding the DR ID of the missing cluster to the `cluster.clusterSpec.dr.excludeClusters` property. The property value is an array of DR IDs. For example, if the DR ID (`cluster.config.deployment.dr.id`) of the lost cluster is "3", you set the property to "{3}":

```
--set cluster.clusterSpec.dr.excludeClusters='{3}'
```

You must set this property for *all* of the clusters remaining in the XDCR environment. If later, you want to add the missing cluster (or another cluster with the same DR ID) back into the XDCR mesh, you will need to reset this property. For example:

```
--set cluster.clusterSpec.dr.excludeClusters=null
```

9.4. Rejoining an XDCR Cluster That Was Previously Removed

If a cluster is removed from the XDCR cluster permanently, by resetting the DR role, or through exclusion by the other clusters, it is still possible to rejoin that cluster to the XDCR network. To do that you must reinitialize the cluster and, if it was forcibly excluded, remove the exclusion from the current members of the network. (Note, the following procedure is *not* necessary if the cluster was removed temporarily by setting the number of replicas to zero.)

First, if the cluster was forcibly removed by exclusion, you must remove the exclusion from the current members of the XDCR network by clearing the `cluster.clusterSpec.dr.excludeClusters` property (removing the missing cluster's ID from the array):

```
--set cluster.clusterSpec.dr.excludeClusters=null
```

Then you must restart the cluster you want to rejoin, reinitializing the cluster's contents with the `cluster.clusterSpec.initForce` property and setting the appropriate properties (such as the DR role and connection properties):

```
--set cluster.clusterSpec.initForce=TRUE  
--set cluster.config.deployment.dr.role="xdcr"  
--set cluster.clusterSpec.replicas=3
```

Once the cluster rejoins the XDCR network and synchronizes with the current members, be sure to reset the `cluster.clusterSpec.initForce` property to false.

Appendix A. Volt Diagnostics Tools

Volt Active Data provides a number of diagnostic tools to assist in managing and troubleshooting VoltDB databases. On Kubernetes, some of these tools are available with additional automation in a separately installable diagnostics tools pod. The diagnostics tools include:

- Data Collection (kcollect) — Collects historical and performance data from the nodes of the cluster, for both the database and the server, then compresses them into a single file on the diagnostics tools pod for easy backup or submission to Volt support as necessary.
- Interactive SQL (ksqlcmd) — Interactive access to the database using the same commands and directives available using **sqlcmd** on bare metal.
- Network performance monitoring (meshmonitor) — Monitors network performance between cluster nodes.
- Snapshots (ksnapshot) — Saves a snapshot of the current database contents, then collects all of the files from the cluster nodes and compresses them into a single file on the diagnostics tools pod.

You can start the diagnostics pod using Helm, then create a terminal session (using **kubectrl exec**) and run the individual tools from the command line. This appendix explains how to start and stop the pod and provides a brief description of the individual tools. See the online help for each tool for further information.

A.1. Starting and Stopping the Diagnostics Pod

You start the diagnostics pod using Helm in the same way you start a VoltDB database. First make sure you have the latest Volt Active Data Helm charts by doing a **helm repo update**:

```
$ helm repo update
```

Next, start the pod with the **helm install** command. The diagnostics pod is independent of the Volt database and the Operator. So it does not matter whether you start the tools pod before or after the database. For example, the following command starts the diagnostics pod using the *volt-diagnostics* chart to create a release called *tools*:

```
$ helm install tools voltdb/volt-diagnostics
```

The diagnostics helm chart does not create a deployment resource, it only creates the temporary pod. This approach has several benefits:

- The pod is consistently named using the release name plus the suffix *-volt-diagnostics*.
- The pod is not automatically rescheduled if Kubernetes stops it, which can be beneficial if your cluster runs low on resources.

When you have finished using the diagnostics pod, you can remove it by simply uninstalling it. For example, the following command deletes the diagnostics pod, *tools*, created in the previous example:

```
$ helm uninstall tools
```

A.1.1. Using the Diagnostics Pod With a Persistent Volume (PV)

The only writable space on the diagnostics pod itself is the `/tmp` directory. If you need more space for collecting cluster data (for example, when using the **collect** command) we encourage you to attach persistent storage. You can attach additional storage in two ways:

- If you already have defined a PersistentVolumeClaim (PVC), set the `diagnostics.pvc.claimName` property pointing to the PVC:

```
$ helm install tools voltdb/volt-diagnostics \
  --set diagnostics.pvc.claimName=tooloutput
```

- If no PVC is defined, but there is a existing Storage Class, you can use the `diagnostics.pvc.storage.className` property and a PVC will be created for you when the diagnostics pod starts:

```
$ helm install tools voltdb/volt-diagnostics \
  --set diagnostics.pvc.storage.className=standard
```

Attaching a PVC to the diagnostics pod provides an additional writeable target for the output of the diagnostics tools, **collect**, **meshmonitor**, and **snapshot**. Use the output qualifiers, either `--outdir` or `--output-file`, to save output to the PVC.

A.1.2. Using the Diagnostic Tools Securely with TLS/SSL

If the cluster being analyzed has security enabled using TLS/SSL, the diagnostics pod must provide the necessary key information to establish connection to the cluster. The tools themselves also connect to VoltDB cluster so must be aware of the TLS/SSL settings. You can configure this information when you start the pod using the `diagnostics.ssl.*` properties. For example:

```
$ helm install tools voltdb/volt-diagnostics \
  --set-file diagnostics.ssl.truststore.file=mytrust.jks \
  --set diagnostics.ssl.truststore.password=mySSLpassword
```

Note that the password is optional. Once TLS is configured for the diagnostics pod, those security settings are used automatically when invoking the diagnostics tools from within the pod.

A.2. Running the Diagnostic Tools

The diagnostics pod comes with four primary tools customized for use with Kubernetes. To use these tools you must first connect to the diagnostics pod using the **kubect** **exec** command. Note that the Kubernetes pod name is the Helm release name plus the suffix `-volt-diagnostics`. For example, the following command accesses the tool pod for the Helm release named `tools`:

```
$ kubectl exec tools-volt-diagnostics -it -- /bin/bash
```

When you connect to the command shell for the diagnostics pod, it displays a welcome message, listing all of the VoltDB databases running in the current namespace. For example:

```
-----
Welcome!
Volt Kubernetes Basic Environment information
-----
```

```
Helm releases installed:
  mydb
Default release used: "mydb"
  There are 6 pods in READY state
  There are 0 pods in NOT READY state
-----
```

From the command prompt, you can issue commands for any of the diagnostics tools. For example:

```
$ ksqlcmd --query="select * from contestants" --output-file=/tmp/results.txt
$ kcollect --release="mydb" --outdir=/tmp
```

A.3. Diagnostic Tools

The following sections give a summary of each of the diagnostics tools. See the online help for detailed instructions on commands and qualifiers.

A.3.1. Collect (kcollect)

The **kcollect** command collects logs, error files, and configuration information for both the base system and the database. This information can help debug issues with the database operation.

```
kcollect [--release={release-name}] [--dir={directory-path}]
```

Additional qualifiers:

```
--context={context-name}
--namespace={namespace-name}
--no-cores
--no-zip
--password={password}
--username={user-name}
```

If there is only one database release running in the namespace, the release name defaults to that release. Otherwise, you should specify the release name of the database cluster with the `--release` qualifier. Similarly, the output directory where the resulting collection is stored defaults to the `/tmp` directory. If you are using an additional PVC, you can use the `--dir` qualifier to redirect the output.

Note

The command for running the collect tool changed in version 1.2.0. See the online help for the **helm voltadmin** command on the diagnostics pod for information on running the collect function on earlier versions.

A.3.2. Mesh Monitor (meshmonitor, V1.5.0 and later)

The **meshmonitor** command starts network monitoring on the database cluster, recording network performance between the cluster nodes. This tool is useful in diagnosing issues such as network delays and instability, mysterious timeouts and hangs, and scheduling problems that impact database performance.

```
meshmonitor [qualifier]... start
```

```
meshmonitor [qualifier]... stop
```

```
meshmonitor [qualifier]... collect
```

Qualifiers:

```
--dir={directory-path}
--release={release-name}
```

There are three commands available for mesh monitor. The **start** command starts monitoring on the cluster nodes. The **stop** command stops monitoring, collects the results, and saves them as a compressed file in the output directory (`/tmp` by default). The **collect** command collects current results and saves them without stopping the monitoring.

For a quick analysis of potential network issues, you can use the **meshmonitor start** and **stop** commands to start monitoring for an hour or so, then stop and collect the data. This is useful if you are seeing problems with database connectivity or latency and want to see if network performance may be the cause.

For extended monitoring, you can use the **meshmonitor start** command to initiate monitoring, then intermittently use the **collect** command to collect and review the results without interrupting further monitoring. When you are done, you can use the **stop** command to cancel monitoring and collect the final data.

Note that monitoring continues only as long as the database pod is running. If the pod or pods stop for any reason (for example a crash or a reboot), the monitoring will not automatically restart. You can use the **meshmonitor start** command to restart monitoring in this situation.

Note

The **meshmonitor** command was introduced in version 1.5.0 of the diagnostic pod. The network monitoring code itself was added to the VoltDB server software in V14.2.0. So, for the time being, monitoring cannot be turned on for earlier versions of VoltDB and the **meshmonitor** command will report that the command is unsupported for the VoltDB version running on the cluster. However, the monitoring code will be included in all future releases of VoltDB and, in upcoming updates, it will be added to the earlier LTS versions as well.

A.3.3. Snapshot (ksnapshot, V1.3.0 and later)

The **ksnapshot** command takes a native database snapshot of the database's current contents, collects the resulting files and saves them as a compressed file in the output directory (`/tmp` by default).

```
ksnapshot [--release={release-name}] [--outdir={directory-path}]
```

If there is only one database release running in the namespace, the release name defaults to that release. Otherwise, you should specify the release name of the database cluster with the `--release` qualifier. Similarly, the output directory where the resulting collection is stored defaults to the `/tmp` directory. If you are using an additional PVC, you can use the `--outdir` qualifier to redirect the output.

Note

The **ksnapshot** command was introduced in version 1.3.0 of the diagnostics pod.

A.3.4. Sqlcmd (ksqlcmd)

The **ksqlcmd** command lets you enter interactive SQL commands and execute stored procedures (including system procedures) on the database.

```
ksqlcmd [--release={release-name}]
```

The **ksqlcmd** command is identical to the **sqlcmd** command except with the `--release` qualifier in place of the `--servers` and `--port` qualifiers for specifying which database you want to connect

to. By default, the command connects to the "default" release listed in the welcome message when you connect to the diagnostics pod. The command also uses the TLS/SSL security configuration specified for the diagnostics pod when connecting to the database, as described in Section A.1.2, “Using the Diagnostic Tools Securely with TLS/SSL”.

Appendix B. VoltDB Helm Properties

You communicate with the VoltDB Operator, and Kubernetes itself, through the Helm charts that VoltDB provides. You can also specify additional Helm properties that customize what the Helm charts do. The properties are hierarchical in nature and can be specified on the Helm command line either as one or more YAML files or as individual arguments. For example, you can specify multiple properties in a YAML file then reference the file as part of your command using the `--values` or `-f` argument, like so:

```
$ helm install mydb voltdb/voltdb --values myoptions.yaml
```

Or you can specify the properties individually in dot notation on the command line using the `--set` flag, like so:

```
$ helm install mydb voltdb/voltdb \
  --set cluster.clusterSpec.replicas=5 \
  --set cluster.config.deployment.cluster.kfactor=2 \
  --set cluster.config.deployment.cluster.sitesperhost=12
```

For arrays and lists, you can specify the values in dot notation by enclosing the list in braces and then quoting the command as required by the shell you are using. For example:

```
$ helm upgrade mydb voltdb/voltdb -reuse-values
  --set cluster.clusterSpec.excludeClusters='{1,3}'
```

In YAML, you specify each element of the property on a separate line, following each parent element with a colon, indenting each level appropriately, and following the last element with the value of the property. On the command line you specify the property with the elements separated by periods and the value following an equals sign. So in the preceding `install` example, the matching YAML file for the command line properties would look like this:

```
cluster:
  clusterSpec:
    replicas: 5
  config:
    deployment:
      cluster:
        kfactor: 2
        sitesperhost: 12
```

Many of the properties have default values; the following tables specify the default values where applicable. You do not need to specify values for all of the properties. In fact, you can start a generic VoltDB database specifying only the license file. Otherwise, you need only specify those properties you want to customize.

Finally, the properties are processed in order and can be overridden. So if you specify different values for the same property in two YAML files and as a command line argument, the latter YAML file setting overrides the first and the command line option overrides them both.

B.1. How to Use the Properties

The following sections detail all of the pertinent Helm properties that you can specify when creating or modifying the VoltDB Operator and its associated cluster. The properties are divided into categories and each category identified by the root elements common to all properties in that category:

- Top-Level Kubernetes Options

- Kubernetes Cluster Startup Options
- Network Options
- VoltDB Database Startup Options
- VoltDB Database Configuration Options
- Operator Configuration Options
- Metrics Configuration Options
- Volt Management Center (VMC) Configuration Options

For the sake of brevity and readability, the properties in the tables are listed by only the unique elements of the property after the root. However, when specifying a property in YAML or on the command line, you must specify all elements of the full property name, including both the root and the unique elements.

B.2. Top-Level Kubernetes Options

The following properties affect how Helm interacts with the Kubernetes infrastructure.

Table B.1. Top-Level Options

Parameter	Description	Default
cluster.enabled	Configure VoltDB Cluster as part of Helm operation (set false if you are installing operator and cluster separately)	true
cluster.serviceAccount.create	If true, create and use service account for VoltDB cluster node containers	true
cluster.serviceAccount.name	If not set and create is true, a name is generated using the full-name template	""

B.3. Kubernetes Cluster Startup Options

The following properties affect the size and structure of the Kubernetes cluster that gets started, as well as the startup attributes of the VoltDB cluster running on those pods.

Table B.2. Options Starting with cluster.clusterSpec...

Parameter	Description	Default
.additionalAnnotations	Additional custom pod annotations	{ }
.additionalLabels	Additional custom pod labels	{ }
.additionalStartArgs	Additional arguments for the 'voldb start' command issued in the pod container	[]
.additionalVolumeMounts	Pod volumes to mount into the container's filesystem; cannot be modified once set	[]
.additionalVolumes	Additional list of volumes that can be mounted by node containers	[]
.affinity	Kubernetes node affinity	{ }
.allowRestartDuringUpdate	Allow VoltDB cluster restarts if necessary to apply user-requested configuration changes. May include automatic save and restore of database.	false

Parameter	Description	Default
.adminOperationTimeout	Timeout for activity check for admin actions such as pause/stop/shutdown. If not specified 120 seconds is used as default. If set less than 120 it will use default 120 seconds.	0
.autoScaling.enabled	Enable/disable auto-scaling. Also used to reset a failed state by disable/enable sequence	false
.autoScaling.maxReplicas	Maximum scale up limit. Effective value will be rounded up to nearest multiple of kfactor+1	16
.autoScaling.maxRetries	Maximum number of times a failed elastic operation will be re-tried. 0 means no retries	0
.autoScaling.metrics.cpu.scaleDown	The threshold that the value of the CPU metric must cross downwards for a cluster scale down	0
.autoScaling.metrics.cpu.scaleUp	The threshold that the value of the CPU metric must cross upwards for a cluster scale up	0
.autoScaling.metrics.cpu	The 'CPU percent usage' metric: the average value of the PERCENT_USED values reported in the CPU statistics	{ }
.autoScaling.metrics.idletime.scaleDown	The threshold that the value of the idle time metric must cross upwards for a cluster scale down	0
.autoScaling.metrics.idletime.scaleUp	The threshold that the value of the idle time metric must cross downwards for a cluster scale up	0
.autoScaling.metrics.idletime	The 'idle time' metric: the average value of the PERCENT values reported in the IDLETIME statistics. Note: lower values require scale up, higher values require scale down	{ }
.autoScaling.metrics.memused.scaleDown	The threshold that the percent integer value of the MEMUSED metric must cross downwards for a cluster scale down	0
.autoScaling.metrics.memused.scaleUp	The threshold that the percent integer value of the MEMUSED metric must cross upwards for a cluster scale up	0
.autoScaling.metrics.memused	The 'memory used' metric: the average PERCENT integer value of the RSS over PHYSICALMEMORY values reported in the MEMORY statistics	{ }
.autoScaling.metrics.rss.scaleDown	The threshold that the value of the RSS metric must cross downwards for a cluster scale down	0
.autoScaling.metrics.rss.scaleUp	The threshold that the value of the RSS metric must cross upwards for a cluster scale up	0
.autoScaling.metrics.rss	The 'resident set size' metric: the average value of the RSS values reported in the MEMORY statistics	{ }
.autoScaling.metrics.tps.scaleDown	The threshold that the value of the TPS metric must cross downwards for a cluster scale down	0
.autoScaling.metrics.tps.scaleUp	The threshold that the value of the TPS metric must cross upwards for a cluster scale up	0
.autoScaling.metrics.tps	The 'transactions per second' metric: the average value of the TPS values reported in the LATENCY statistics	{ }
.autoScaling.metrics	Lists the thresholds for the monitored metrics, indexed by metric name: cpu, idletime, rss, tps	{ }

Parameter	Description	Default
.autoScaling.minReplicas	Minimum scale down limit. Effective value will be rounded up to nearest multiple of kfactor+1	kfactor + 1
.autoScaling.notificationInterval	The duration, in seconds, between notification events reporting that an elastic operation is ongoing. 0 means no notification	0
.autoScaling.retryTimeout	Defines the duration, in seconds, to wait for a retried operation to start. If the timeout expires and the operation didn't start, auto-scaling will be stopped.	60
.autoScaling.stabilizationWindow .scaleDown	The duration, in seconds, that a 'scaleDown threshold crossed' condition must remain true in order to trigger an elastic remove operation	300
.autoScaling.stabilizationWindow .scaleUp	The duration, in seconds, that a 'scaleUp threshold crossed' condition must remain true in order to trigger an elastic add operation	300
.clusterInit .classesConfigMapRefName	Name of pre-created Kubernetes configmap containing stored procedure classes	""
.clusterInit.licenseSecretRefName	Name of pre-created Kubernetes secret containing Volt license, using key 'license.xml'	""
.clusterInit.logConfigMapName	Name of pre-created Kubernetes config map containing custom logging configuration, using key 'log4j2.properties' or 'log4j.xml'	""
.clusterInit .schemaConfigMapRefName	Name of pre-created Kubernetes configmap containing schema configuration	""
.customEnv	Key-value map of additional environment variables to set in all VoltDB node containers	{ }
.disableFinalizers	Disables Helm finalizers to permit cluster deletion. WARNING: many resources will require manual cleanup.	false
.deletePVC	Delete and cleanup generated PVCs when VoltDBCluster is deleted, requires finalizers to be enabled (on by default)	false
.dr.excludeClusters	User-specified list of clusters not part of XDCR	[]
.dr.forceDrop	Indicate if you want to drop cluster from XDCR without producer drain.	false
.elasticRemove.checkInterval	Time in seconds to wait between checks of the status of an ongoing elastic remove operation. A value of 10 seconds or more is recommended to let other workflows be executed by the operator	10
.elasticRemove.ignore	Can be set to disabled_export to force ignoring the disabled exports, since elastic remove waits for all exports to be drained prior to removing the nodes	``
.elasticRemove.restart	Requests the restart of an elastic remove operation currently in the FAILED state. Value must be nonzero and also different from the last value used for restart	0
.elasticRemove.shutdownDelay	Specifies the number of minutes to wait before shutting down the nodes being removed. Must be greater than 0 if topics are being used, otherwise the elastic remove fails	0
.elasticRemove.update	Requests an update of the parameters of an ongoing elastic remove operation, e.g. ignore or shutdownDelay. Value must be nonzero and also different from the last value used for update	0

Parameter	Description	Default
.elasticReset	Requests a reset of the elastic remove information in the cluster status. Reserved for VoltDB support. Value must be nonzero and also different from the last value used for reset	0
.enableInServiceUpgrade	Enable rolling upgrade of software version rather than requiring full cluster restart (V13.1.0 or later).	false
.env.VOLTDDB_GC_OPTS	VoltDB cluster java runtime garbage collector options (VOLTDDB_GC_OPTS)	""
.env.VOLTDDB_HEAPCOMMIT	Commit VoltDB cluster heap at startup, true/false (VOLTDDB_HEAPCOMMIT)	""
.env.VOLTDDB_HEAPMAX	VoltDB cluster heap size, integer number of megabytes (VOLTDDB_HEAPMAX)	""
.env.VOLTDDB_OPTS	VoltDB cluster additional java runtime options (VOLTDDB_OPTS)	""
.env.VOLTDDB_REGION_LABEL_NAME	Override for region label on node	""
.env.VOLTDDB_ZONE_LABEL_NAME	Override for zone label on node	""
.forceStopNode	Enable or disable force stop node (V12.2 or later)	false
.image.pullPolicy	Image pull policy	Always
.image.registry	Image registry	docker.io
.image.repository	Image repository	voltldb/voltldb-enterprise
.image.tag	Image tag	Same as global.voltldbVersion
.inServiceUpgrade.delay	FOR TESTING PURPOSES ONLY: Specifies the delay in seconds upgrading pods to new image.	false
.initForce	Always init --force on VoltDB node start/restart. WARNING: This will destroy VoltDB data on PVCs except snapshots.	false
.livenessProbe.enabled	Enable/disable livenessProbe; see Kubernetes documentation for probe settings	true
.maintenanceMode	VoltDB Cluster maintenance mode (pause all nodes)	false
.maxPodUnavailable	Maximum pods allowed to be unavailable in Pod Disruption Budget	kfactor
.network	Per-connection IP advertising for multi-NIC pods (e.g. Multus); see values.yaml for structure/examples. Requires VoltDB 13.3.16+, 14.3.6+, or 15.3.0+.	{ }
.nodeSelector	Node labels for pod assignment	{ }
.persistentVolume.hostpath.enabled	Use HostPath volume for local storage of VoltDB. This node storage is often ephemeral and will not use PVC storage classes if enabled.	false

Parameter	Description	Default
.persistentVolume.hostpath.path	HostPath mount point.	"/ data/volt- db/"
.persistentVolume.size	Persistent Volume size per pod (VoltDB Node)	32Gi
.persistentVolume.storageClassName	Storage Class name to use, otherwise use default	""
.podSecurityContext	Pod security context defined by Kubernetes	See file val- ues.yaml
.podTerminationGracePeriodSeconds	Duration in seconds the pod needs to terminate gracefully.	30
.priorityClassName	Pod priority defined by an existing PriorityClass	""
.readinessProbe.enabled	Enable/disable readinessProbe; see Kubernetes documentation for probe settings	true
.replicas	Pod (VoltDB Node) replica count; scaling to 0 will shutdown the cluster gracefully	3
.resources	CPU/Memory resource requests/limits	{ }
.safeMode	Start VoltDB in 'safe recovery' mode (see documentation)	false
.securityContext	Container security context defined by Kubernetes	See file val- ues.yaml
.statefulSetUpdateMode	Override default decision on applying stateful set changes; not normally set. Values are restart, stop, terminate.	``
.ssl.certificateFile	PEM-encoded certificate chain used by the operator to verify VoltDB cert when TLS/SSL is enabled	""
.ssl.clientCertFile	PEM-encoded private key and certificate identifying operator when mutual TLS is required	""
.ssl.insecure	If true, skip VoltDB certificate verification by the operator when TLS/SSL is enabled	false
.startupProbe.enabled	Enable/disable startupProbe; see Kubernetes documentation for probe settings	true
.stoppedNodes	User-specified list of stopped VoltDB nodes, by pod ordinal (0, 1, ...)	[]
.storageConfigs	Optional storage configs for provisioning additional persistent volume claims automatically	[]
.takeSnapshotOnShutdown	Takes a snapshot when cluster is shut down by scaling to 0. One of: NoCommandLogging, Always, Never. NoCommandLogging means 'only if command logging is disabled'.	"NoCom- mandLog- ging"
.tempVolumeSizeLimit	Limit on temporary storage volume mounted on /tmp; decimal number optionally followed by one of M/Mi/G/Gi/T/Ti	no limit
.tolerations	Pod tolerations for node assignment (see Kubernetes documen- tation)	[]
.topologySpreadConstraints	Describes how a group of pods ought to spread across topology (see Kubernetes documentation)	[]
.useCloudNativePlacementGroup	Enable or disable cloud native placement group in VoltDB	false

B.4. Network Options

The following properties specify what ports to use and the port-mapping protocol.

Table B.3. Options Starting with `cluster.serviceSpec...`

Parameter	Description	Default
<code>.adminNodePort</code>	Port to expose VoltDB Admin service on each node, type NodePort only	31211
<code>.adminPortEnabled</code>	Enable exposing admin port with the VoltDB service	true
<code>.adminPort</code>	Admin port number to be exposed on VoltDB service	21211
<code>.clientNodePort</code>	Port to expose VoltDB Client service on each node, type NodePort only	31212
<code>.clientPortEnabled</code>	Enable exposing client port with the VoltDB service	true
<code>.clientPort</code>	Client port number to be exposed on VoltDB service	21212
<code>.dr.annotations</code>	Additional custom service annotations	{ }
<code>.dr.enabled</code>	Create single DR service for DR	false
<code>.dr.externalTrafficPolicy</code>	VoltDB DR service external traffic policy	""
<code>.dr.ipFamilies</code>	Select IPv4/IPv6 protocols used by DR replication service	[]
<code>.dr.override</code>	Allows per-pod-service overrides of serviceSpec	[]
<code>.dr.publicIPFromService</code>	Operator will wait to get the public IP address from the service status set by Kubernetes	false
<code>.dr.replicationNodePort</code>	Kubernetes service ports[].nodePort for VoltDB replication service on each node, type NodePort only. If -1 is specified, kubernetes will select a random unused port	31555
<code>.dr.replicationPort</code>	Kubernetes service ports[].port for the VoltDB DR replication service	5555
<code>.dr.staticIP</code>	Single static IP for DR service use when creating LoadBalancers single DR service	``
<code>.dr.type</code>	VoltDB DR service type, valid options are ClusterIP (default), LoadBalancer, or NodePort	""
<code>.externalIPs</code>	List of IP addresses at which the VoltDB service is available	[]
<code>.externalTrafficPolicy</code>	VoltDB service external traffic policy (options Cluster, Local)	Cluster
<code>.http.sessionAffinityConfig.clientIP.timeoutSeconds</code>	Timeout override for http.sessionAffinity=ClientIP (onboard http only)	10800
<code>.http.sessionAffinity</code>	SessionAffinity override for the onboard HTTP service	ClientIP
<code>.kafka.annotations</code>	Additional custom Service annotations	{ }
<code>.kafka.availableIPs[]</code>	Available IPs and IP-ranges to use when creating LoadBalancers on a per-pod basis	[]
<code>.kafka.externalTrafficPolicy</code>	Kafka service external traffic policy	""
<code>.kafka.publicIPFromService</code>	Operator will wait to get the public IP address from the service status set by Kubernetes	false
<code>.kafka.topicsNodePort</code>	NodePort for the kafka service when the service type is NodePort	31092

Parameter	Description	Default
.kafka.topicsPort	Kafka replication exposed Service port	9092
.kafka.type	Kafka service type, valid options are ClusterIP (default), LoadBalancer, or NodePort	""
.loadBalancerIP	VoltDB Load Balancer IP	""
.loadBalancerSourceRanges	VoltDB Load Balancer Source Ranges	[]
.perpod.dr.enabled	Enable DR services on a per-pod basis	false
.perpod.dr.externalTrafficPolicy	VoltDB DR service external traffic policy for per pod DR services.	""
.perpod.dr.ipFamilies	Select IPv4/IPv6 protocols used by per pod DR replication services	[]
.perpod.dr.replicationPort	Kubernetes service ports[].port for the perpod VoltDB DR replication services	5555
.perpod.dr.startReplicationNodePort	Starting Kubernetes service ports[].nodePort for perpod VoltDB replication service, type NodePort only. Sequential allocation from start. -1 means kubernetes will select a random unused port	32555
.perpod.metrics.enabled	Enables metrics k8s service for each pod	false
.perpod.publicIPFromService	Operator will wait to get the public IP address from the service status set by Kubernetes	false
.perpod.staticIPs[]	Available IPs and IP-ranges to use when creating LoadBalancers on a per-pod basis	[]
.perpod.type	VoltDB service type, valid options are ClusterIP (default), LoadBalancer, or NodePort	""
.service.metrics.type	Metrics port service type (options ClusterIP, NodePort, and LoadBalancer)	ClusterIP
.type	VoltDB service type (options ClusterIP, NodePort, and LoadBalancer)	ClusterIP
.vmcNodePort	Port to expose Volt Management Center service on each node, type NodePort only (onboard httpd only)	31080
.vmcPort	Volt Management Center web interface Service port (onboard httpd only)	8080
.vmcSecureNodePort	Port to expose Volt Management Center secure service on each node, type NodePort only (onboard httpd only)	31443
.vmcSecurePort	Volt Management Center secure web interface Service port (onboard httpd only)	8443

B.5. VoltDB Database Startup Options

The following properties affect how Helm interacts with the VoltDB cluster and specific initialization options, such as the initial schema and procedure classes.

Table B.4. Options Starting with cluster.config...

Parameter	Description	Default
.auth.credSecretName	Name of the premade secret containing Operator admin username and password. This overrides auth.username and auth.password values and avoids including the password in yaml.	""
.auth.password	Operator admin password used to access VoltDB; required if security is enabled. Superseded by credSecretName when provided.	""
.auth.username	Operator admin user name used to access VoltDB; required if security is enabled. Superseded by credSecretName when provided.	voltodb-operator
.classes	Map of optional jar files containing stored procedures	{ }
.licenseXMLFile	VoltDB Enterprise license file	{ }
.log4jcfgFile	Custom Log4j configuration file	{ }
.schemas	Map of optional schema files containing data definition statements	{ }
.usersSecretName	Name of pre-created Kubernetes secret containing users configuration, using key 'users.yaml'. If defined, no users should be specified in 'cluster.config.deployment'.	""

B.6. VoltDB Database Configuration Options

The following properties define the VoltDB database configuration.

Table B.5. Options Starting with cluster.config.deployment...

Parameter	Description	Default
.avro.namespace	Avro namespace (since V11)	""
.avro.prefix	Avro configuration prefix	""
.avro.properties	Avro configuration properties	{ }
.avro.registry	Avro registry URL	""
.cluster.kfactor	K-factor to use for database durability and data safety replication	1
.cluster.sitesperhost	SitesPerHost for VoltDB Cluster	8
.commandlog.enabled	Command logging for database durability (recommended)	true
.commandlog.frequency.time	How often the command log is written, by time (milliseconds)	200
.commandlog.frequency.transactions	How often the command log is written, by transaction commands	2147483647
.commandlog.logsize	Command logging allocated disk space (MB)	1024
.commandlog.synchronous	Transactions do not complete until logged to disk	false
.compoundproc.callsperstage	Maximum number of procedure invocations per stage of each compound procedure	10
.compoundproc.queueulimit	Maximum number of compound procedures queued to the compound procedure executor	10000

Parameter	Description	Default
.compoundproc.threads	Number of threads used by the compound procedure executor	Max(availableProcessors, 4)
.dr.conflictretention	Automatic pruning of XDCR conflict logs; integer followed by one of m/h/d, for minutes/hours/days(since V11)	""
.dr.connection.connectiontimeout	Connection timeout for connecting to producer cluster in milliseconds.	""
.dr.connection.enabled	Specifies whether this DR connection is enabled	false
.dr.connection.preferredSource	Cluster ID of preferred DR source	""
.dr.connection.receiveTimeout	Timeout receiving response from producer cluster in milliseconds	""
.dr.connection.source	If role is replica or xdc: list of host names or IP addresses of remote nodes	""
.dr.connection.ssl.secretName	Optional pre-made secret containing truststore data, including password if needed	""
.dr.connection.ssl.truststore.file	Optional truststore file used to verify the identity of the remote VoltDB cluster; defaults to truststore of this cluster, unless sslSecret is set	""
.dr.connection.ssl.truststore.password	Password for truststore file specified in truststore.file	""
.dr.connection.ssl-host-check	Enable validation of Subject Alternative Name in received certificates	false
.dr.consumerlimit.maxbuffers	Enable DR consumer flow control; either maxsize or maxbuffers must be specified	""
.dr.consumerlimit.maxsize	Enable DR consumer flow control; either maxsize or maxbuffers must be specified. maxsize is an integer optionally followed by 'g' (gigabytes) or 'm' (megabytes)	""
.dr.exportconflicts	Optional setting. If false, XDCR conflict logs are not generated into export stream and therefore are not exported. Conflicts are still handled and shown in statistics.	true
.dr.id	Unique cluster id, 0-127	0
.dr.role	Role for this cluster, either xdc or none	xdc
.dr.schemachange.enabled	Enable DR consumer to continue while compatible schema changes are being made (since V12)	false
.dr.schemachange.truncate	Enable values to be truncated if a VARCHAR column is wider on another cluster while schema changes are being made	false
.export.configurations	List of export configurations (complex structure)	[]
.export.configurations.enabled	Is this export enabled?	false
.export.configurations.exportconnectorclass	Class name for custom exporters	""
.export.configurations.properties	List of type-specific properties	[]
.export.configurations.reconnect	Reconnect interval, integer followed by 's' or 'm' (V12 or later)	"120s"
.export.configurations.target	Name of export target	""
.export.configurations.threadpool	Name of thread pool to use for export processing	""

Parameter	Description	Default
.export.configurations.type	Export type (file, custom, etc.)	"file"
.export.defaultpoolsize	Size of the default export thread pool (if not set, sites per host value used); only used for Volt V13+	Not set
.heartbeat.timeout	Internal VoltDB cluster verification of presence of other nodes (seconds)	90
.httpd.enabled	Determines if HTTP API daemon is enabled (external VMC service is recommended instead)	false
.httpd.jsonapi.enabled	Determines if JSON over HTTP API is enabled; requires http.enabled	false
.httpd.port	Specifies port for HTTP; default varies according to whether TLS/SSL enabled	8080 or 8443
.import.configurations	List of import configurations (complex structure)	[]
.import.configurations.enabled	Is this import enabled?	"false"
.import.configurations.format	Format of import data	""
.import.configurations.formatProperties	List of formatter-specific properties (only used for Volt V13+)	[]
.import.configurations.module	Name of jar file for custom importer	""
.import.configurations.nickname	Arbitrary but unique name for import configuration (Volt V14+)	"import-NNNNNNNN"
.import.configurations.priority	Priority of transactions for this import configuration	4
.import.configurations.properties	List of type-specific properties	[]
.import.configurations.type	Import type e.g. kafka	""
.import.configurations.version	The kafka version (only used for Volt V13+)	"10"
.metrics.enabled	Enables cloud-native metrics system on each VoltDB pod. It is an alternative to using Prometheus Agent.	false
.metrics.interval	How often the metrics system prepares a new metrics slice to be sourced by an external system like prometheus. Supported unit indicators are s (seconds), m (minutes), h (hours), d (days).	60s
.metrics.maxbufferize	Limits memory that the metric system can use for internal metric buffering (system will always retain at least one metrics slice in the buffer). In megabytes.	16
.partitiondetection.enabled	Controls detection of network partitioning. Setting this false is deprecated in V14, and doing so will be ignored in V15+.	true
.paths.commandlog.path	Directory path for command log	/pvc/voltdb/voltdbroot/command_log
.paths.commandlogsnapshot.path	Directory path for command log snapshot	/pvc/voltdb/voltdbroot/command_log_snapshot
.paths.droverflow.path	Directory path for disaster recovery overflow	/pvc/voltdb/voltdbroot/

Parameter	Description	Default
		dr_overflow
.paths.exportcursor.path	Directory path for export cursors	/pvc/voltdb/voltdbroot/export_cursor
.paths.exportoverflow.path	Directory path for export overflow	/pvc/voltdb/voltdbroot/export_overflow
.paths.largequeryswap.path	Directory path for large query swapping	/pvc/voltdb/voltdbroot/large_query_swap
.paths.snapshots.path	Directory path for snapshots. Must be located in a writeable directory, not the root (as init --force will rename existing snapshot folder).	/pvc/voltdb/voltdbroot/snapshots
.security.enabled	Controls whether user-based authentication and authorization are used	false
.security.ldap.group	List of LDAP groups and their mapping to VoltDB roles	[]
.security.ldap.group.name	Distinguished name of LDAP group	""
.security.ldap.group.nickname	Arbitrary but unique name for this group entry (Volt V14+)	name/role
.security.ldap.group.role	Comma-list of Volt roles equivalent to this LDAP group	""
.security.ldap.groupclass	Name of the LDAP schema's objectClass defining a group of users	"groupOfUniqueNames"
.security.ldap.groupmemberid	Name of the LDAP schema's objectClass defining a group of users	"uniqueMember"
.security.ldap.password	Password corresponding to LDAP server username, required	""
.security.ldap.rootdn	Distinguished Name of the root of the LDAP schema that defines users and groups, required	""
.security.ldap.server	URL for LDAP server, required; as 'ldap://server:port', or 'ldaps://server:port', port optional	""
.security.ldap.ssl.truststore.file	Truststore file used to validate LDAPS server certificate (Java KeyStore format)	""
.security.ldap.ssl.truststore.password	Password for LDAP truststore file	""
.security.ldap.timeout	Timeout, in seconds, for requests to the LDAP server	10
.security.ldap.user	Username used by VoltDB for read-only access on LDAP server, required	""
.security.ldap.userclass	Name of the LDAP schema's objectClass containing user information	"inetOrgPerson"
.security.ldap.useruid	Name of the LDAP attribute in the userObjectClass that should contain the username provided by the VoltDB client	"uid"

Parameter	Description	Default
.security.provider	Sets authentication provider as hash (local) or ldap (using a customer-specified LDAP/LDAP server)	hash
.snapshot.enabled	Enable/disable periodic automatic snapshots	true
.snapshot.frequency	Frequency of automatic snapshots (in s,m,h)	24h
.snapshot.prefix	Unique prefix for snapshot files	AUTOS-NAP
.snapshot.retain	Number of snapshots to retain	2
.snmp.authkey	SNMPv3 authentication key if protocol is not NoAuth	voltbdbauthkey
.snmp.authprotocol	SNMPv3 authentication protocol. One of: SHA, MD5, NoAuth	SHA
.snmp.community	Name of SNMP community	public
.snmp.enabled	Enables or disables use of SNMP	false
.snmp.privacykey	SNMPv3 privacy key if protocol is not NoPriv	voltbdbprivacykey
.snmp.privacyprotocol	SNMPv3 privacy protocol. One of: AES, DES, 3DES, AES192, AES256, NoPriv	AES
.snmp.target	Host name or IP address, and optional port (default 162), for SNMP server	""
.snmp.username	Username for SNMPv3 authentication; else SNMPv2c is used	""
.ssl.clientauthrequired	Enables mTLS on external port; client applications must be configured with a certificate	false
.ssl.crl.configName	Name of pre-created Kubernetes configmap containing CRL files (overrides the explicit CRL files if defined)	""
.ssl.crl.file	Map of optional explicit CRL files (e.g. --set-file cluster.config.deployment.ssl.crl.file.1=crl1, --set-file cluster.config.deployment.ssl.crl.file.2=crl2, etc...)	{ }
.ssl.dr	Extends TLS/SSL security to the DR port (5555).	false
.ssl.drclientauthrequired	Enables mTLS on DR port; DR consumer will use local cluster keystore to identify itself to remote	false
.ssl.enabled	Enable or disable TLS/SSL on the cluster. Other properties control activation of TLS/SSL for specific ports and features.	false
.ssl.external	Extends TLS/SSL security to all external ports (default admin 21211, client 21212).	false
.ssl.internal	Extends TLS/SSL security to the internal port (default 3021).	false
.ssl.keystore.file	Keystore file for cluster, in Java Key Store format; ignored if sslSecret is used	""
.ssl.keystore.password	Password for VoltDB keystore file	""
.ssl.sslSecret.certSecretName	Pre-made secret containing keystore and truststore data, optionally including passwords	""
.ssl.sslSecret.passwordSecretName	Pre-made secret containing password for keystore/truststore, if password is not in the secret named by certSecretName	""

Parameter	Description	Default
.ssl.truststore.file	Truststore file, in Java Key Store format; ignored if sslSecret is used	""
.ssl.truststore.password	Password for VoltDB truststore file	""
.systemsettings.clockskew.interval	Interval of the scheduled clock skew collection (minutes). 0 is allowed and it disables collection. Interval cannot be less than 0 and if set below such value it will be reset to default.	60
.systemsettings.compaction.interval	Interval to indicate how often memory compaction should run (seconds)	60
.systemsettings.compaction.maxcount	Set a target block count compaction should try and achieve if there is memory fragmentation	1
.systemsettings.elastic.duration	Target value for the length of time each rebalance transaction will take (milliseconds)	50
.systemsettings.elastic.throughput	Target value for rate of data processing by rebalance transactions (MB)	2
.systemsettings.flushinterval.dr.interval	Interval for flushing DR data (milliseconds)	1000
.systemsettings.flushinterval.export.interval	Interval for flushing export data (milliseconds)	4000
.systemsettings.flushinterval.minimum	Interval between checking for need to flush (milliseconds)	1000
.systemsettings.priorities.batchsize	Modifies priority scheduling algorithm to execute multiple requests before rescheduling	25
.systemsettings.priorities.dr.priority	Priority for DR requests (1-8, 1 is highest priority)	5
.systemsettings.priorities.enabled	Enables priority scheduling of requests by VoltDB cluster (true/false)	false
.systemsettings.priorities.maxwait	Modifies priority scheduling by setting a limit on time waiting while higher priority requests execute (millisecs)	1000
.systemsettings.priorities.snapshot.priority	Priority for snapshot requests (1-8, 1 is highest priority)	6
.systemsettings.procedure.copyparameters	If set, mutable array parameters should be copied before processing	true
.systemsettings.procedure.loginfo	Threshold for long-running task detection (milliseconds)	10000
.systemsettings.query.timeout	Timeout on SQL queries (milliseconds)	10000
.systemsettings.resourcemonitor.disklimit.commandlog.alert	Alert level for disk use (in GB or as percentage, empty is unlimited)	""
.systemsettings.resourcemonitor.disklimit.commandlog.size	Limit on disk use (in GB or as percentage, empty is unlimited)	""
.systemsettings.resourcemonitor.disklimit.commandlogsnapshot.alert	Alert level for disk use (in GB or as percentage, empty is unlimited)	""
.systemsettings.resourcemonitor.disklimit.commandlogsnapshot.size	Limit on disk use (in GB or as percentage, empty is unlimited)	""
.systemsettings.resourcemonitor.disklimit.droverflow.alert	Alert level for disk use (in GB or as percentage, empty is unlimited)	""

Parameter	Description	Default
.systemsettings.resourcemonitor.disklimit.droverflow.size	Limit on disk use (in GB or as percentage, empty is unlimited)	""
.systemsettings.resourcemonitor.disklimit.exportoverflow.alert	Alert level for disk use (in GB or as percentage, empty is unlimited)	""
.systemsettings.resourcemonitor.disklimit.exportoverflow.size	Limit on disk use (in GB or as percentage, empty is unlimited)	""
.systemsettings.resourcemonitor.disklimit.snapshots.alert	Alert level for disk use (in GB or as percentage, empty is unlimited)	""
.systemsettings.resourcemonitor.disklimit.snapshots.size	Limit on disk use (in GB or as percentage, empty is unlimited)	""
.systemsettings.resourcemonitor.disklimit.topicsdata.alert	Alert level for disk use (in GB or as percentage, empty is unlimited)	""
.systemsettings.resourcemonitor.disklimit.topicsdata.size	Limit on disk use (in GB or as percentage, empty is unlimited)	""
.systemsettings.resourcemonitor.frequency	Interval between Resource Monitor resource checks (seconds)	60
.systemsettings.resourcemonitor.memorylimit.alert	Alert level for memory use (in GB or as percentage)	70%
.systemsettings.resourcemonitor.memorylimit.size	Limit on memory use (in GB or as percentage)	80%
.systemsettings.resourcemonitor.memorylimit.compact	Enable memory compaction	true
.systemsettings.snapshot.priority	Delay factor for snapshot work; larger values cause longer waits between bursts of execution. Ignored if systemsettings.priorities.enabled is true.	6
.systemsettings.snapshot.autotune.enabled	Autotuning for snapshot work (for beta test only, in 14.1.0)	false
.systemsettings.temptables.maxsize	Limit the size of temporary database tables (MB)	100
.task.mininterval	Task scheduling parameter	0
.task.maxfrequency	Task scheduling parameter	0.0
.task.threadpools	Threadpools available for task execution	[]
.task.threadpools.host.size	Threadpool host size	0
.task.threadpools.partition.size	Threadpool partition size	0
.threadpools	List of thread pools to create	[]
.threadpools.pool.name	Thread pool name	""
.threadpools.pool.size	Thread pool size	""
.topics.broker.properties	Kafka topics: broker configuration properties	[]
.topics.broker	Kafka topics: broker configuration	``
.topics.enabled	Kafka topics: enabled or not (supported since V11)	true
.topics.threadpool	Kafka topics: threadpool to use	``
.topics.topic.allow	List of roles allowed to access the topic	``
.topics.topic.format	Format of topic message	``

Parameter	Description	Default
.topics.topic.name	Topic name	``
.topics.topic.opaque	Is this an opaque topic?	false
.topics.topic.priority	Priority for topics requests (if priority scheduling is enabled)	4
.topics.topic.procedure	Procedure to invoke upon getting message	``
.topics.topic.properties	Topic configuration properties	[]
.topics.topic.retention	Topic retention policy	``
.topics.topic	List of topics	[]
.users.expires	Expiration date for this user account	None
.users.name	A VoltDB user name	""
.users.password	User password	""
.users.plaintext	If false, password was pre-hashed	false
.users.roles	Comma-list of roles assigned to user	""
.users	List of VoltDB users to be added to the deployment	[]

B.7. Operator Configuration Options

The following properties configure the Volt Operator, which is in turn responsible for managing the startup and operation of all other Volt components.

Table B.6. Options Starting with operator...

Parameter	Description	Default
.affinity	Kubernetes node affinity (node affinity and pod anti-affinity)	{ }
.cleanupCustomResource	Attempt to delete CRD when uninstalling the Helm chart	false
.cleanupNamespaceClusters	Delete ALL VoltDB clusters in the namespace when the operator Helm chart is deleted	false
.cleanupImagePullPolicy	Pull policy for operator image used in cleanup mode	IfNotPresent
.debug.enabled	Debug level logging	false
.enabled	Create VoltDB Operator to manage clusters	true
.image.pullPolicy	Image pull policy	Always
.image.registry	Image registry	docker.io
.image.repository	Image repository	voltodb/volt-db-operator
.image.tag	Image tag	Same as Chart version
.livenessProbe	Standard probe; see Kubernetes documentation for probe settings	See file values.yaml
.logformat	Log encoding format for the operator (console or json)	json
.nodeSelector	Node labels for pod assignment	{ }
.podAnnotations	Additional custom pod annotations	{ }

Parameter	Description	Default
.podLabels	Additional custom pod labels	{ }
.readinessProbe	Standard probe; see Kubernetes documentation for probe settings	See file values.yaml
.replicas	Pod replica count	1
.resources	CPU/Memory resource requests/limits	See file values.yaml
.securityContext	Kubernetes container security context	See file values.yaml
.serviceAccount.create	If true, create & use service account for VoltDB operator containers	true
.serviceAccount.name	If not set and create is true, a name is generated using the full-name template	""
.startupProbe	Standard probe; see Kubernetes documentation for probe settings	See file values.yaml
.tolerations	Kubernetes pod tolerations for node assignment	[]
.topologySpreadConstraints	Kubernetes topology spread constraint	{ }

B.8. Metrics Configuration Options

Properties starting with `metrics...` were used to configure the standalone VoltDB Prometheus agent. However, the Prometheus agent has been deprecated in favor of per pod metrics. See Section 6.1, “Using Prometheus to Monitor VoltDB” for more information on using the current metrics system.

B.9. Volt Management Center (VMC) Configuration Options

The following properties start and configure the web-based Volt Management Center auxiliary service.

Table B.7. Options Starting with vmc...

Parameter	Description	Default
.affinity	Kubernetes node affinity	{ }
.enabled	Create VoltDB VMC Service	true
.image.pullPolicy	Image pull policy	Always
.image.registry	Image registry	docker.io
.image.repository	Image repository	voltldb/volt-db-vmc
.image.tag	Image tag	From global.voltldb-Version
.nodeSelector	Node labels for pod assignment	{ }
.resources	CPU/Memory resource requests/limits	See file values.yaml

Parameter	Description	Default
.service.ssl.enabled	Enable TLS/SSL for the VMC service: true/false/auto. 'Auto' uses cluster settings.	"auto"
.service.ssl.keystore.file	Contents of the keystore file for the VMC service (if not using auto or sslSecret)	""
.service.ssl.keystore.password	Password for the keystore file for the VMC service	""
.service.ssl.sslSecret.certSecretName	Pre-made secret containing truststore data, optionally including password	""
.service.ssl.sslSecret.passwordSecretName	Pre-made secret containing password for truststore, if not in secret certSecretName	""
.service.ssl.truststore.file	Contents of the trust store file for the VMC service (if not using auto or sslSecret)	""
.service.ssl.truststore.password	Password for the trust store file for the VMC service, if used	""
.serviceSpec.externalTrafficPolicy	External Traffic Policy of VMC Service	See file values.yaml
.serviceSpec.type	Type of VMC Service	See file values.yaml
.sessionLogging	Enable logging of client HTTP sessions and connections to VoltDB (V13.3.3+, V14.1.0 and later)	false
.tempVolumeSizeLimit	Limit on temporary storage volume mounted on /tmp; see cluster .clusterSpec.tempVolumeSizeLimit	same as cluster
.tolerations	Pod tolerations for node assignment	[]
.topologySpreadConstraints	Topology Spread Constraint	{ }
.voltdb.adminPort	VoltDB API admin port, uses 21211 if not specified	21211
.voltdb.credSecretName	Name of the premade secret.	""
.voltdb.port	VoltDB API port, uses 21212 if not specified	21212
.voltdb.ssl.enabled	Enable TLS/SSL for VMC to VoltDB traffic: true/false/auto. 'Auto' uses cluster settings if needed.	"auto"
.voltdb.ssl.keystore.file	Contents of key store file if VoltDB requires mutual auth, V14.2+ (if not using auto or sslSecret)	""
.voltdb.ssl.keystore.password	Password for key store file, if used	""
.voltdb.ssl.sslSecret.certSecretName	Pre-made secret containing truststore data, optionally including password	""
.voltdb.ssl.sslSecret.passwordSecretName	Pre-made secret containing password for truststore, if not in secret certSecretName	""
.voltdb.ssl.truststore.file	Contents of the VoltDB trust store file (if not using auto or sslSecret)	""
.voltdb.ssl.truststore.password	Password for the VoltDB trust store file, if used	""