



ZENworks® 2020 Update 1

Vertica Reference Guide

June 2020

Legal Notice

© Copyright 2008 - 2020 Micro Focus or one of its affiliates.

The only warranties for products and services of Micro Focus and its affiliates and licensors ("Micro Focus") are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Micro Focus shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

Contents

About This Guide	5
1 Overview	7
ZENworks and Vertica	7
Points to Consider before moving to Vertica	8
2 Components required for Vertica	9
ZooKeeper	9
Apache Kafka	9
Kafka Change Data Capture Workflow	10
3 Prerequisites	13
System Requirements and Other Recommendations	13
4 Configuring Vertica	15
Workflow to configure Vertica in the Zone	15
Procedure to Configure Vertica and Kafka on a Single Server	16
Enabling Kafka	19
Enabling Vertica	21
Managing Vertica Memory Requirements	22
Migrating Data to Vertica	23
5 Managing Vertica, Kafka and ZooKeeper	25
Adding Servers to the Vertica, Kafka and ZooKeeper Cluster	25
Prerequisites	25
ZooKeeper	26
Kafka	26
Vertica	28
Maintaining ZooKeeper, Vertica and Kafka Clusters	29
Maintaining the ZooKeeper Cluster	29
Maintaining the Kafka Cluster	30
Maintaining the Vertica Cluster	32
Monitoring the Status of Vertica, Kafka and ZooKeeper	32
Viewing the Configuration Status	32
Monitoring Diagnostics	33
Connecting to Vertica	34
Stopping and Starting Vertica Services	35
Shutting Down a Vertica Node	36
6 Disaster Preparedness and Recovery	37
Backing up and Restoring the Vertica Database	37

Creating Configuration File for Vertica Full Local Backup	38
Creating a Full Local Backup	38
Restoring from a backup	39
Changing the ZooKeeper role from the first Primary Server to the second Primary Server.	39
7 Debugging and Log Location	41
ZooKeeper	41
Configuration Issues	41
General Debugging	41
Kafka	42
Configuration Issues	42
General Debugging	42
Vertica	43
Configuration Issues	43
Vertica Management Console.	43
Data Migration	44
Firewall Issues	44
8 Removing Vertica	47
Disabling and Removing Vertica from the Zone	47
Enabling Vertica	48
9 Troubleshooting	51
ZooKeeper	51
Kafka	51
Vertica	53
A Appendix 1: Mounting a New Hard Disk for Vertica Catalog	55
Glossary	57
B Documentation Updates	59
October 2020: ZENworks 2020 Update 1	59

About This Guide

This *ZENworks Vertica Reference Guide* includes information to help you successfully enable and configure Vertica and its related components.

Audience

This guide is intended for ZENworks administrators.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation.

Additional Documentation

ZENworks is supported by other documentation (in both PDF and HTML formats) that you can use to learn about and implement the product. For additional documentation, see the [ZENworks documentation Web site](#).

1 Overview

Vertica is an analytics database platform that is specifically built for Big Data analytics. Vertica is built for organizations that generate large volumes of data frequently. Compared to traditional relational database systems (RDBMS), Vertica provides faster query performance and scalability while analyzing these large volumes of data. The most important features of Vertica include:

- ♦ Being a columnar database, it stores information in columns and only reads the columns required to answer a query.
- ♦ Highly scalable as Vertica runs on a cluster of commodity hardware. As and when the volume of data generated increases, you can add more nodes to the Vertica cluster to increase processing speeds.
- ♦ Compresses data size to a significant extent, which improves analytic performance while processing historical or trending data.
- ♦ Queries and loads data on a real-time basis up to 10X faster than traditional RDBMS.
- ♦ Efficiently handles multiple operations with multiple concurrent users.

Some other features include:

- ♦ Supports standard programming interfaces that include ODBC, JDBC, ADO.NET, and OLEDB.
- ♦ Uses database designer and administration that allows you to control Vertica with minimal administration effort.
- ♦ Analyzes semi-structured data along with structured data.

For more information on the Vertica database, see the documentation available at www.vertica.com.

ZENworks and Vertica

In ZENworks, Vertica is not meant to replace the existing ZENworks primary database but is supposed to work in conjunction with the existing database.

You can leverage the capabilities of Vertica when using the Dashboards feature in ZENworks. While all dashlets, with the exception of trending charts, will be displayed when connected to the existing RDBMS; to reduce any delays that you might face while querying large volumes of data in the dashlets, you can configure the Vertica database.

As Vertica stores and reads information from columns leading to faster processing speeds, the data in the dashlets will be processed and displayed much faster compared to an RDBMS database. Also, as Vertica has the ability to compress data size, it can store historical data, thereby enabling you to view historical trends related to patches and CVEs. For more information on the Dashboards feature in ZENworks, see the [ZENworks 2020 Dashboard Reference](#) guide.

Being an in-house Micro Focus product, the Vertica database is license free for up to 20 TB of data, beyond which you will have to purchase a license for Vertica.

Vertica can be configured only on an Appliance server. To configure the Vertica database, and to migrate and obtain data from Vertica, you also need the ZooKeeper and Apache Kafka components to be enabled and running at all times. When you upgrade to or install ZENworks 2020, ZooKeeper is by default enabled on the first Primary Server. However, Apache Kafka like Vertica needs to be enabled on an Appliance server using the ZMAN utility .

Points to Consider before moving to Vertica

This section is intended to help you make an informed decision about moving to the Vertica database. Some important aspects that you need to carefully consider before enabling Vertica in your zone are:

- ♦ ZENworks does not mandate the use of Vertica. However, you should configure Vertica if:
 - ♦ You have more than 10 thousand devices in the zone and if you use the Common Vulnerabilities and Exposure (CVE) feature. In such a scenario, as significant amount of data is stored in the database, it is mandatory to have the Vertica database installed in the zone.

NOTE: For zones with a lesser number of devices, you can continue using the existing RDBMS, as all features (except for trend charts) will continue to work seamlessly.

- ♦ You want to view the trending charts displayed within the Patch Tracker and CVE Tracker dashlets. Irrespective of the number of devices in the zone, trend data will be displayed only when Vertica is enabled.
- ♦ When you configure Vertica in the zone:
 - ♦ Vertica will only process data for the Dashboard feature.
 - ♦ ZENworks will process data from Vertica only for the Security dashlets. The bundle dashlets will read data from the existing RDBMS unless you manually configure these dashlets (Device Assignment Status dashlet and the Bundle Deployment Status dashlet) to obtain data from Vertica, which is documented in the [Software Distribution Guide](#).

NOTE: The bundle dashlets will work seamlessly with the existing RDBMS, even when you have a sizable number of devices in the zone.

- ♦ ZENworks will continue to use the existing RDBMS as the primary database.
 - ♦ The RDBMS will continue to receive data from the ZENworks Primary Server. As an initial step, data will be migrated in bulk from the RDBMS to Vertica, after which the RDBMS will continuously sync data with Vertica at an interval of every 10 to 15 minutes, using the Apache Kafka and ZooKeeper components.
 - ♦ The RDBMS will retain all data, except for the following which will be directly stored in Vertica:
 - ♦ CVE trend data and Patch trend data: If you remove Vertica from the zone, you will lose this data.
 - ♦ CVE Vulnerability Status data.
- ♦ ZENworks currently does not provide you with a control to disable Vertica automatically, if you no longer want to use it in your zone. You need to manually disable it in the zone, which is documented in [Removing Vertica](#).

2 Components required for Vertica

The Vertica database is dependent on the Apache Kafka and ZooKeeper components and these components should be up and running at all times for Vertica to function effectively. This section provides a brief overview of these components. The topics covered in this section are as follows:

- ♦ [“ZooKeeper” on page 9](#)
- ♦ [“Apache Kafka” on page 9](#)

ZooKeeper

ZooKeeper acts as a coordination service that provides flexible and robust synchronization within distributed systems. ZooKeeper forms an integral part of Kafka as it keeps track of the status of the Kafka cluster nodes and the Kafka topics and partitions. The leader election between the Kafka brokers is also performed by using ZooKeeper when the leader node fails.

Apache ZooKeeper is already packaged with the ZENworks 2020 build. When the build is deployed on the first Primary Server, ZooKeeper is automatically enabled in the zone. However, you can expand the ZooKeeper cluster as per your scalability requirements and to provide fault tolerance.

NOTE: Besides Vertica, other services in ZENworks require the ZooKeeper services to be up and running at all times. Therefore, even if you do not have Vertica configured in your zone, ZooKeeper should always be active. You also need to ensure that the firewall allows client connections from other Primary Servers like those in the DMZ, to the ZooKeeper services on port 6789. For more information, see [Firewall Issues](#).

For more information on adding additional servers to the ZooKeeper cluster, see [Managing Vertica, Kafka and ZooKeeper](#).

For more information on monitoring the ZooKeeper services, see [Monitoring Diagnostics](#).

NOTE: If vulnerabilities are reported while running the Qualys scan on a Primary Server that has ZooKeeper configured on it, then to resolve the issue follow the procedure documented in [TID 7024819](#).

Apache Kafka

Apache Kafka is a distributed publish-subscribe messaging system that enables passing of messages from one system to another, while handling large volumes of data. Kafka can be enabled on a single Appliance Primary Server or can be run as a cluster on one or more Appliance servers that can span multiple data centers. Each server in the cluster is called a broker. Kafka is run as a cluster to ensure high availability of its services by replicating Kafka topics or messages to multiple Kafka brokers. Kafka requires ZooKeeper to co-ordinate between the servers within the Kafka cluster.


Apache Kafka is pre-packaged in the ZENworks 2020 virtual appliance build but will be in a disabled state. Using the ZMAN utility, you need to enable Kafka and its related services on an Appliance server. For more information on enabling Kafka, see [Enabling Kafka](#).

In the Vertica data migration workflow, Apache Kafka is required to sync data with Vertica, after the bulk data migration process. When Vertica is enabled, the ZENworks server will continue to send data to the existing RDBMS. After bulk data migration, if any data is added or modified in the RDBMS, Kafka syncs these changes with Vertica at an interval of every 10 to 15 minutes.

For more information on Kafka, see [Kafka Reference Guide](#).

Kafka Change Data Capture Workflow

The following diagram is a graphical representation of the Kafka workflow:



A description of each component in this architecture is as follows:

- ♦ **Kafka Cluster:** A group of servers or nodes that are connected to each other to achieve a common objective is a cluster. Each of the servers or nodes in the cluster, will have one instance of Kafka broker running.
- ♦ **Kafka Broker:** One or more servers that are added in a Kafka cluster are called brokers.
- ♦ **Apache ZooKeeper:** Kafka uses ZooKeeper to manage and co-ordinate between brokers. It notifies Kafka Producers and Consumers of any broker failures.
- ♦ **Kafka Producers:** The processes that publish messages to Kafka brokers. In this case, Kafka connectors are created as soon as Kafka is enabled. These connectors are created for each table in the RDBMS database and is responsible for identifying changes in these tables and publishing them to Kafka.
- ♦ **Kafka Consumers:** A service that reads data from Kafka brokers. In this case, the ZENworks loader stream processors subscribes to data from Kafka and passes this information to the Vertica database.

The Kafka workflow is as follows:

1. The Kafka connectors (Kafka producer) identifies changes in the respective RDBMS database tables.
2. These changes are published to a topic within a Kafka broker. Kafka maintains messages as categories and these categories are called topics. To modify the interval at which the Kafka connector identifies changes in the RDBMS tables and publishes this data to Kafka, you can

update the `connector-configs.xml` file (available at `etc/opt/novell/zenworks`) and run the `zman srkcen` command to re-configure the connectors. For more information on this command, see [Maintaining Kafka Connect](#).

The topics are further broken down into partitions for speed and scalability. These partitions are replicated across multiple brokers, which is used for fault tolerance. The replicas are created based on the replication count specified while enabling the Kafka cluster. Each message within a partition is maintained in a sequential order, which is identified as an offset, also known as a position.

3. The ZENworks loader stream processors (consumers) subscribe to a topic, Kafka offers the current offset of the topic to the consumer and saves the offset in the ZooKeeper cluster.
4. The ZENworks loader stream processor receives the message and processes it to the Vertica database and the Kafka broker receives an acknowledgment of the processed message.

Kafka also uses a Schema Registry, which is a separate entity that producers and consumers talk to, for sending and retrieving schemas that describe the data models for the messages.

For more information on the Kafka Architecture, see the [Confluent](#) docs.

3 Prerequisites

A few general prerequisites that need to be met before enabling and using Vertica are as follows:

- ♦ **Appliance Server:** ZENworks 2020 virtual appliance build is deployed in the zone. Only Appliance servers can be a part of the Vertica and Kafka clusters.
- ♦ **Zone Infrastructure Rights:** The **Modify Zone Infrastructure** rights should be configured for your administrator. For more information on enabling these rights, see [ZENworks Administrator Accounts and Rights Reference](#).
- ♦ **Two-way SSL Communication:** To secure the Kafka cluster, Kafka allows clients to connect over SSL and uses the same Primary Server certificate on which it is installed. However, if the server certificate is issued by an external CA, then for Kafka to work seamlessly ensure that both Client Authentication and Server Authentication is configured in the Extended Key Usage (EKU) attribute.
- ♦ **TCP and UDP Ports:** Ensure that the default ports required to communicate with the Kafka, ZooKeeper and Vertica services are open at all times. For more information on these default ports, see [ZENworks 2020 TCP and UDP Ports](#).

System Requirements and Other Recommendations

The minimum hardware requirements and other recommendations for the Vertica and Kafka servers (nodes) are:

- ♦ **Single Node Cluster:** A single node (server) cluster is recommended if you have less than 10000 devices in the zone. For a single node cluster, you can configure both Kafka and Vertica on a single node.
- ♦ **Multiple Node Cluster:** If you have more than 10000 devices in your zone, then you require more than one Vertica node. As the minimum size of a cluster is three nodes, a minimum three node cluster is required for such zones.

For multiple node cluster, it is recommended that Vertica, Kafka and the Collection role are configured on different servers. As Vertica and Kafka are Disk I/O (input/output) intensive, Kafka should be enabled on a server that does not have Disk I/O intensive roles configured on it. As the Collection role is also Disk I/O intensive, Kafka should not be configured on the same server on which the Collection role is enabled.

- ♦ **Hardware Configuration:** The minimum hardware requirements for the Vertica and Kafka servers are available at the location https://www.novell.com/documentation/zenworks-2020/pdfdoc/vertica_system_requirements/vertica_system_requirements.pdf. The minimum hardware configuration was derived based on the number of devices in the zone and the number of administrators accessing the Dashboards simultaneously.

Ensure that the hardware configuration for all the servers within the Vertica cluster are the same.

- ♦ **Network Configuration:** Ensure that all the nodes in the cluster are in the subnet to enable internal cluster communication. For more information on the network configuration recommendations, see the [Vertica](#) documentation.

- ♦ **Server Hostname:** While adding servers or brokers to the Vertica, Kafka or ZooKeeper clusters, ensure that the servers do not have multiple hostnames.
- ♦ **Memory Allocation:** To ensure appropriate memory is allocated for the Vertica and Kafka components and the ZENworks services on the server, run the **Calibrate Memory Configure action** on the server on which either Vertica and Kafka are installed. This action should be run after Vertica and Kafka are installed and before you proceed with bulk data migration. For more information, see [Managing Vertica Memory Requirements](#).

For more information on other recommendations, see the [Vertica](#) documentation.

4 Configuring Vertica

The topics covered in this section are as follows:


- ♦ [“Workflow to configure Vertica in the Zone” on page 15](#)
- ♦ [“Procedure to Configure Vertica and Kafka on a Single Server” on page 16](#)
- ♦ [“Enabling Kafka” on page 19](#)
- ♦ [“Enabling Vertica” on page 21](#)
- ♦ [“Managing Vertica Memory Requirements” on page 22](#)
- ♦ [“Migrating Data to Vertica” on page 23](#)

Workflow to configure Vertica in the Zone

The following workflow provides a high-level overview of configuring the Vertica database and its related components on a single server zone.

NOTE: You do not have to enable ZooKeeper in your zone, as it is by default enabled on your first Primary Server as soon as you upgrade or install the ZENworks 2020 release version.

- 1. Enable Kafka in the zone:** Packaged with the ZENworks 2020 virtual appliance build, Apache Kafka will initially be in a disabled state. Use the ZMAN utility to enable Kafka and its related services (Kafka connect and schema registry). For more information, see [Procedure to Configure Vertica and Kafka on a Single Server](#).
- 2. Enable Vertica in the zone:** Packaged with the ZENworks 2020 virtual appliance build, Vertica will initially be in a disabled state. You need to use the ZMAN utility to enable Vertica in the zone. Ensure that you enable Kafka in the zone, before enabling Vertica. For more information, see [Procedure to Configure Vertica and Kafka on a Single Server](#).
- 3. Migrate data:** Migration of data from the existing RDBMS to Vertica includes two parts, which is represented in the following illustration:



Step 1 - Bulk Data Migration: As soon as Kafka and Vertica are enabled in the zone, you need to perform bulk data migration from the existing RDBMS to Vertica by executing a Configure action. After migrating data in bulk, ZENworks will start processing queries from Vertica for all the dashlets in ZENworks. For more information, see [Procedure to Configure Vertica and Kafka on a Single Server](#).

Step 2 - Change Data Capture: After the initial bulk migration process, as and when data is modified in the existing RDBMS, for these changes to reflect in Vertica, the RDBMS syncs with Vertica at an interval of every 15 minutes. Apache Kafka along with ZooKeeper is responsible for performing this sync action. For more information, see [Procedure to Configure Vertica and Kafka on a Single Server](#).

Procedure to Configure Vertica and Kafka on a Single Server

A summary of the steps to be followed to configure Vertica and Kafka on a single server are provided in this section. For more information on each of the steps, click the links specified within each step. To test the feature, it is recommended that you follow this procedure first before adding more servers to the Kafka and Vertica clusters. Both Kafka and Vertica can be enabled on the same server or on two different servers.

Prerequisites

Before enabling Vertica and Kafka in the zone, you need to make a note of the following requirements:


Procedure

NOTE: ZooKeeper should by default be enabled on your first Primary Server. If you want to add more servers to the ZooKeeper cluster, then see [Adding Servers to the Vertica, Kafka and ZooKeeper Cluster](#). To verify whether the ZooKeeper service is up and running, navigate to ZCC and monitor the status in the Diagnostics page.

1 Enable Kafka at the Cluster Level and then at the Broker level

1a At Cluster Level: Execute the following ZMAN command and specify the replication count.

```
zman server-role-kafka-configure-cluster (zman srkcc)
```




```
kish-janus:~ # zman server-role-kafka-configure-cluster -c 1
Username:Administrator
Password:*****
Cluster configured successfully.
kish-janus:~ #
```

For more information, see [Enabling Kafka](#).

1b At Broker Level: Execute the following ZMAN command and specify the GUID, DNS or path of the Appliance server on which Kafka should be enabled.

```
zman server-role-kafka-add-broker (zman srkab)
```



```
kish-janus:~ # zman server-role-kafka-add-broker --servers=kish-
Username:Administrator
Password:*****
All prerequisite validations have passed.
Adding new broker(s) to kafka cluster, this may take a few minutes.
Proceeding with queue operations for the new broker(s).
The operation on server kish- has succeeded.
Rebalancing topic partitions..
Partitions for topics successfully rebalanced.
Command completed.
kish-janus:~ #
```

When this command is executed for a server, the Kafka, Kafka connect and Schema Registry services are enabled and started on the specified server.

For more information, see [Enabling Kafka](#).

If you are unable to add a Kafka broker to the cluster, then check the logs on the server that is to be added to the cluster as detailed in the section [Kafka](#).

2 View the status of Kafka configuration in ZCC, by navigating to **Configuration > Performance Upgrade**. To continuously monitor the status of Kafka and its components, navigate to the Diagnostics page in ZCC. For more information, see [Monitoring the Status of Vertica, Kafka and ZooKeeper](#).

- 3 Enable Vertica in the zone by executing the following command and specify the GUID, DNS, or path of the Appliance server, and the K factor.

```
zman server-role-vertica-create-cluster (zman srvcc)
```

```
kish-janus@ # zman server-role-vertica-create-cluster --servers= -k=0
Username:Administrator
Password:*****

Preparing the servers...

The operation on server kish- has succeeded.

Preparation completed successfully on all the servers.

Starting Vertica installation. This might take a few minutes...For details, see /opt/vertica/log/install.log on kish-

Re-configuring the database.This might take a few minutes...For details, see /opt/vertica/log/adminTools.log.
Database creation or reconfiguration completed successfully.
Post database creation tasks completed successfully.
Updating the K-safety value.
K-safety configured successfully.
Creating the schema.
Adding the Vertica role.
Updating the configuration files.
Configuration files created successfully.
Vertica cluster installed successfully.
The operation on server kish- has succeeded.
Updating configuration files on other primary servers. This might take a few minutes...

The operation on server kish- has succeeded.

The creation of the configuration files has completed successfully.

Vertica cluster created successfully.

Restart the ZENServer and ZENLoader services on all the Primary Servers to re-initialize the ZENworks datamodel.
```

For more information, see [Enabling Vertica](#).

If Vertica cluster configuration fails, then check the logs on the server on which Vertica is to be enabled as detailed in the section [Vertica](#).

- 4 Restart the ZENserver and ZENloader services on all the Primary Servers.
- 5 View the status of Vertica configuration in ZCC, by navigating to **Configuration > Performance Upgrade**. To continuously monitor the status of Vertica and its nodes, navigate to the Diagnostics page in ZCC. For more information, see [Monitoring the Status of Vertica, Kafka and ZooKeeper](#).
- 6 Run the Configure action to calibrate memory requirements for the Vertica and Kafka components. To run the action:

- 6a In the command line utility of the Appliance server on which Vertica and Kafka are enabled, execute the following command:

```
novell-zenworks-configure -c CalibrateMemoryConfigureAction
```

- 6b A message is displayed informing the user that the ZENworks services will automatically restart on the server after the Configure action is run. If you want to proceed with the memory calibration, then press enter to specify the default value as 1 or else enter the value 2.

- 6c Login as a dbadmin user by running the following command: `su dbadmin`

- 6d Run the following commands to restart the Vertica database:

```
1. $ /opt/vertica/bin/admintools -t stop_db -d zenworks -p
<password> -F
```

```
2. $ /opt/vertica/bin/admintools -t start_db -d zenworks -p  
<password>
```

You can obtain the password by running the command `zman server-role-vertica-get-credentials (zman srvgc)` and you need to use the database superuser password to stop and start the Vertica database.

For more information, see [Managing Vertica Memory Requirements](#).

- 7 Export data in bulk to the configured tables in Vertica from the existing RDBMS, by running the following Configure action on one of the servers in which Vertica is enabled:

```
novell-zenworks-configure -c VerticaDBMigrate -Doption=start
```

For more information, see [Migrating Data to Vertica](#).

After bulk data migration, the Change Data Capture process automatically syncs modified or newly added data from the RDBMS to Vertica at an interval of every 10 to 15 minutes.

You can now proceed to add more servers to the ZooKeeper, Kafka, or Vertica clusters. For more information, see [Adding Servers to the Vertica, Kafka and ZooKeeper Cluster](#).

Enabling Kafka

Apache Kafka is already packaged in your ZENworks 2020 Appliance build and will initially be in a disabled state. To enable Kafka, you need to execute ZMAN commands in the command line utility at both the cluster level and broker level. You need to first configure the Kafka cluster settings and subsequently add a broker to it.

Cluster Level:

- ♦ **Configure Kafka Cluster:** Execute the following command:

```
zman server-role-kafka-configure-cluster (zman srkcc)
```

While running this command, you need to specify the replication count. If the values for the remaining parameters are not specified, then the default values are considered. The parameters are:

- ♦ **-a --logRetentionPeriod:** Specify the duration (in hours) for which the Kafka logs should be retained.
- ♦ **-c --replication count:** This ensures that messages remain available when a server in the cluster fails. Specify the number of copies to be maintained for each topic. For a single node cluster, this count can be set as 1.
- ♦ **-l --logRetentionBytes:** Specify the maximum permissible size of the log, beyond which, the existing data is overwritten with the new data. By default the log size is unlimited.
- ♦ **-t --zkSessionTimeout:** Specify the ZooKeeper session timeout (in milliseconds), which is the maximum time that the server waits to establish a connection to ZooKeeper. If the server fails to signal a heartbeat to ZooKeeper within this specified time period, then the server is considered to be dead. A heartbeat request helps identify if the server is still connected to the Kafka cluster. The default value is 30000 milliseconds.
- ♦ **-r --retainDetectedLogsDuration:** Specify the maximum time to retain deleted logs. The default value is 86400000 milliseconds (1 day).

- ♦ **-p --logCleanupPolicy:** Specify the default cleanup policy for segments that exceed the maximum permissible retention window. The possible values are Delete and Compact. The default value is Delete. The Delete policy will remove old segments when the retention time or size limit has reached. The Compact policy will enable log compaction on the topic, which ensures that Kafka will always retain at least the last known value for each message key within the log of data for a single topic partition.
- ♦ **-s --schemaregistryport:** Specify the port on which the Schema Registry is running. The default value is 8081.
- ♦ **-k --kafkaport:** Specify the port on which Kafka listens. The default value is 9093.
- ♦ **-x --connectport:** Specify the port on which Kafka connect listens. The default value is 8083.

For example, `zman server-role-kafka-configure-cluster -c=1`

Broker Level

NOTE: Before adding a Kafka broker, you need to ensure that the following prerequisites are met:

- ♦ The Kafka cluster should already be configured.
 - ♦ The broker or server that is to be added to the cluster should be an Appliance server and should not already be added to the cluster.
 - ♦ The server should be accessible and should have a specific host name. If a server has multiple host names, then the command to add the Kafka broker might fail.
 - ♦ As Kafka requires client authentication to be enabled in the certificate, if you are using an external CA certificate, ensure that client authentication is enabled in the Extended Key Usage (EKU) parameter of the server certificate.
-

- ♦ **Add Kafka brokers:** Execute the following command:

```
zman server-role-kafka-add-broker (zman srkab)
```

This command adds brokers or servers to the configured Kafka cluster. The parameters to be specified are:

- ♦ **--servers:** Specify the appliance server on which Kafka should be enabled. Specify the DNS, GUID or path of the server object (server, server folder or server group) relative to / Devices/Servers.
- ♦ **-i --ignorewarning message (optional):** As Kafka requires client authentication to be enabled in the certificate, if you are using an external CA certificate, ensure that client authentication is enabled and execute the command again with the option `i` set as true. Option `i` enables you to ignore the warning message and proceed with the execution of the command.

For example: `zman server-role-kafka-add-broker --servers=server1.microfocus.com`

When this command is executed for a server, the Kafka, Kafka connect and Schema Registry services are enabled and started on the specified server.

For more information on debugging Kafka configuration issues, see [Debugging and Log Location](#).

You can view the status of the Kafka Cluster configuration by navigating to [Configuration > Performance Upgrade](#). For more information on monitoring the Kafka cluster, see [Monitoring the Status of Vertica, Kafka and ZooKeeper](#).

Enabling Vertica

NOTE: Before enabling Vertica, ensure that all the services within the server are up and running to avoid any port conflicts.

This section details the procedure to enable Vertica. The Vertica database is already packaged with the ZENworks 2020 virtual appliance build and will be in a disabled state. To enable Vertica in the zone, you need to execute the following command on an Appliance server:

```
zman server-role-vertica-create-cluster (zman srvcc)
```

This command creates a Vertica cluster and adds the server to the cluster, based on the specified parameters. The parameters to be specified while executing this command are:

- ♦ **--servers:** Specify the appliance server on which you want Vertica to be installed. Specify the DNS, GUID or path of the server object (server, server folder or server group) relative to / Devices/Servers.
- ♦ **K factor:** Specify the number of replicas of the data that should exist in the database. The K safety sets the fault tolerance for the Vertica cluster, where K can be set as 0 or 1. You need to specify a value based on the number of nodes in the cluster, which is measured as 2k+1. To enable Vertica on a single node cluster, the K factor should be specified as 0.

If you are planning to create a multi-node cluster, then ensure that you update the K factor to 1 or more. For more information on updating the K factor, see [Updating the K Factor](#).

For example: `zman server-role-vertica-create-cluster --servers=server2.microfocus.com -k=0`

When this command is executed, Vertica is installed on the server and the ZENworks database and its related schema are created.

IMPORTANT: After Vertica is enabled in the zone, ensure that you restart the ZENserver and ZENloader services on all the Primary Servers.

For more information on debugging Vertica configuration issues, see [Debugging and Log Location](#).

You can view the status of the Vertica Cluster configuration by navigating to [Configuration > Performance Upgrade](#). For more information on monitoring the Vertica cluster, see [Monitoring the Status of Vertica, Kafka and ZooKeeper](#).

NOTE: After enabling Vertica, if you want to set a specific language option to be used for error messages and instructions, collating sequences, date formats, and so forth, then see the [Vertica](#) documentation. For more information on locales refer to the following links:

- ♦ <https://www.vertica.com/docs/9.2.x/HTML/Content/Authoring/AdministratorsGuide/Diagnostics/Locales/AboutLocales.htm>

- ♦ <https://www.vertica.com/docs/9.2.x/HTML/Content/Authoring/AdministratorsGuide/Diagnostics/Locales/LocaleUTF8Support.htm>
- ♦ <https://www.vertica.com/docs/9.2.x/HTML/Content/Authoring/AdministratorsGuide/Diagnostics/Locales/SupportedLocales.htm>

After changing the Language option, ensure that you are restart the Vertica database. For more information see, [Stopping and Starting Vertica Services](#).

Managing Vertica Memory Requirements

As per the system requirements, the Appliance server on which Vertica is installed, should have a minimum RAM size of 16 GB. Based on the RAM size, you need to allocate appropriate memory for Vertica, Kafka, the RDBMS (if embedded PostgreSQL is installed) and ZENworks services to function effectively on the Appliance server. ZENworks provides you with a Configure action that will automatically calibrate appropriate memory for these components.

This Configure action needs to be executed after installing Vertica and Kafka and before proceeding with bulk data migration. Subsequently, if you want to add more servers to the Vertica or the Kafka clusters, then you need to run this Configure action on each of the servers. The scenarios in which this action should be run are:

- ♦ Server on which Kafka and Vertica are installed
- ♦ Server on which only Vertica is installed
- ♦ Server on which embedded PostgreSQL and Vertica are installed
- ♦ Server on which embedded PostgreSQL and Kafka are installed
- ♦ Server on which only Kafka is installed
- ♦ Server on which embedded PostgreSQL, Kafka and Vertica are installed

NOTE: The Configure action will fail to run, if the server does not meet the minimum memory requirement of 16 GB.

To run this Configure action:

1. In the command line utility of the Appliance server on which Vertica and Kafka are installed, execute the command:

```
novell-zenworks-configure -c CalibrateMemoryConfigureAction
```
2. A message is displayed informing the user that the ZENworks services will automatically restart on the server after the Configure action is run. If you want to proceed with the memory calibration, then press enter to specify the default value as 1 or else enter the value 2.
3. Login as a dbadmin user by running the following command: `su dbadmin`
4. Run the following commands to restart the Vertica database in the order listed below:

```
$ /opt/vertica/bin/admintools -t stop_db -d zenworks -p <password> -F
```

```
$ /opt/vertica/bin/admintools -t start_db -d zenworks -p <password>
```

You can obtain the password by running the `zman` command `zman server-role-vertica-get-credentials (zman srvgc)`. You need to use the database superuser password to stop and start the Vertica database.

The newly configured values for each component are displayed. If you want to view the configured values at a later point in time, then run the following Configure action:

```
novell-zenworks-configure -c DisplayMemoryConfigureAction
```

Migrating Data to Vertica

NOTE: Before bulk data migration, ensure that you run the Configure action to allocate appropriate memory for ZENworks services and for the Vertica and Kafka components. For more information, see [“Managing Vertica Memory Requirements” on page 22](#).

The data migration process to Vertica involves two steps:

- ♦ **Bulk Data Migration:** This process involves exporting of data in bulk to the configured tables in Vertica from the existing RDBMS. To initiate this process, from the command line utility of the Primary Server in which Vertica is installed, run the following Configure action:

```
novell-zenworks-configure -c VerticaDBMigrate -Doption=start
```

If a Configure action that was executed previously had failed, then the start option will resume the operation from the point at which it had failed. To restart the migration process from the beginning, use the Force option. This option will remove existing data from the Vertica database.

For example, `novell-zenworks-configure -c VerticaDBMigrate -Doption=force`

If the migration process fails, then check the logs as detailed in the section [Data Migration](#).

- ♦ **Change Data Capture:** This event is automatically triggered after the bulk migration process. This process syncs modified or newly added data from the RDBMS to Vertica. The sync action is performed at an interval of every 10 to 15 minutes. For the syncing process, a queue action is initiated that creates Kafka connectors and topics for each table that is already migrated. The Kafka connectors identifies updated or modified tables in the RDBMS database and publishes this data to the appropriate topics. The ZENworks loader stream services (Kafka consumer) then subscribes to this data and sends it to the Vertica database.

To monitor the status of the Data Sync process, navigate to the Diagnostics page in ZCC. For more information, see [Monitoring Diagnostics](#).

5 Managing Vertica, Kafka and ZooKeeper

The topics addressed in this section are as follows:

- ♦ [“Adding Servers to the Vertica, Kafka and ZooKeeper Cluster” on page 25](#)
- ♦ [“Maintaining ZooKeeper, Vertica and Kafka Clusters” on page 29](#)
- ♦ [“Monitoring the Status of Vertica, Kafka and ZooKeeper” on page 32](#)
- ♦ [“Connecting to Vertica” on page 34](#)
- ♦ [“Stopping and Starting Vertica Services” on page 35](#)
- ♦ [“Shutting Down a Vertica Node” on page 36](#)

Adding Servers to the Vertica, Kafka and ZooKeeper Cluster

This section details the procedure to move Vertica and its components from a single node cluster to a multiple node cluster.

NOTE: After adding Kafka and Vertica servers to their respective clusters, ensure that you run the Configure action to allocate appropriate memory for these services, on each of the servers added in the cluster. For more information, see [Managing Vertica Memory Requirements](#).

Prerequisites

The prerequisites to be followed before adding servers to their individual clusters, are:

- ♦ Ensure that the hardware and software configuration of the servers that are to be added are the same as the configuration of the existing server within that cluster.
- ♦ Although there is no limit to the number of servers to be added in a cluster, ZENworks recommends a cluster size of 3 servers for ZooKeeper, Kafka and Vertica.
- ♦ For Kafka, before adding servers, the Kafka cluster should already be configured. For more information, see [Procedure to Configure Vertica and Kafka on a Single Server](#).
- ♦ For Vertica, before adding servers, the Vertica cluster should already be enabled. For more information, see [Procedure to Configure Vertica and Kafka on a Single Server](#).
- ♦ The broker or server that is to be added to the Kafka or Vertica cluster should be an Appliance server and should not already be added to the cluster.
- ♦ The server should be accessible and should have a specific host name.

ZooKeeper

Apache ZooKeeper is already packaged with the ZENworks 2020 build. When the build is deployed on the first Primary Server, ZooKeeper is automatically enabled in the zone. However, you can expand the ZooKeeper cluster as per your scalability requirements and to provide fault tolerance. It is recommended to have an odd number of servers in the cluster. A three node cluster still continues to function if one server goes down since the majority servers are still up. A four node cluster also can only survive the loss of a single node since if two nodes go down, the cluster loses its majority. Therefore, no additional fault tolerance that can be achieved by configuring the cluster with an even number of servers.

Adding Servers to the ZooKeeper Cluster

- ♦ Execute the following command to add nodes or servers to the cluster:

```
zman server-role-zookeeper-add-server (zman srzas)
```

The parameter to be specified is:

- ♦ **--servers:** Specify a comma separated list of Primary Servers on which ZooKeeper should be enabled. You can specify the DNS, GUID or path of the server object (server, server folder or server group) relative to /Devices/Servers.

For example: `zman server-role-zookeeper-add-server --servers=server2.microfocus.com,server3.microfocus.com,server4.microfocus.com,server5.microfocus.com`

For more information on debugging ZooKeeper configuration issues, see [ZooKeeper](#).

Kafka

A summary of the steps to be followed to add brokers to the Kafka cluster are as follows:

NOTE: As Kafka requires client authentication to be enabled in the certificate, if you are using an external CA certificate, ensure that client authentication is enabled.

1. Execute the command `zman server-role-kafka-add-broker (zman srkab)` to add brokers or servers to the configured Kafka cluster. Specify a comma separated list of the server DNS, GUID or path of the server object (server, server folder or server group) relative to /Devices/Servers.

For example: `zman server-role-kafka-add-broker --servers=server1.microfocus.com,server2.microfocus.com`

2. Execute the command `zman server-role-kafka-update-cluster (zman srkuc)` to modify the replication count.

For example: `zman server-role-kafka-update-cluster -c=3`

Adding Servers

NOTE: The server that you are adding should be accessible and should have a specific host name. If a server has multiple host names, then the command to add the Kafka broker will fail.

- ♦ Execute the following command to add brokers or servers to the configured Kafka cluster.

```
zman server-role-kafka-add-broker (zman srkab)
```

The parameters to be specified are:

- ♦ **--servers:** Specify a comma separated list of appliance servers on which Kafka should be enabled. You can specify the DNS, GUID or path of the server object (server, server folder or server group) relative to /Devices/Servers.
- ♦ **-i --ignorewarning message (optional):** As Kafka requires client authentication to be enabled in the certificate, if you are using an external CA certificate, ensure that client authentication is enabled and execute the command again with the option `i` set as true. Option `i` enables you to ignore the warning message and proceed with the execution of the command.

For example: `zman server-role-kafka-add-broker --servers=server1.microfocus.com,server2.microfocus.com`

IMPORTANT: When a broker is being added to the zone, tasks to remove or re-configure servers cannot be performed on other servers.

When this command is executed for a server, the Kafka, Kafka connect and Schema Registry services are enabled and started on the specified server. The server is added to the cluster. When a broker is added in the cluster, ZENworks re-configures other existing brokers in the cluster. When re-configuration fails in one these brokers, you need to run the `zman` command (`zman srkrcb`) to re-configure the broker. For more information on this command, see [Maintaining the Kafka Cluster](#).

Updating Replication Count

- ♦ Execute the following command to modify the existing cluster parameters such as the replication count.

```
zman server-role-kafka-update-cluster (zman srkuc)
```

The parameters are:

- ♦ **-c --replication count:** This ensures that messages remain available when a server in the cluster fails. Specify the number of copies to be maintained for each topic. For a three node cluster, this count can be set as 3, which indicates that a replica of the topic will be created on each server.
- ♦ **-l --logRetentionBytes:** Specify the maximum permissible size of the log, beyond which, the existing data is overwritten with the new data. By default the log size is unlimited.

- ♦ **-t --zkSessionTimeout:** Specify the ZooKeeper session timeout (in milliseconds), which is the maximum time that the server waits to establish a connection to ZooKeeper. If the server fails to signal a heartbeat to ZooKeeper within this specified time period, then the server is considered to be dead. A heartbeat request helps identify if the server is still connected to the Kafka cluster. The default value is 30000 milliseconds.
- ♦ **-r --retainDetectedLogsDuration:** Specify the maximum time to retain deleted logs. The default value is 86400000 milliseconds (1 day).
- ♦ **-p --logCleanupPolicy:** Specify the default cleanup policy for segments that exceed the maximum permissible retention window. The possible values are Delete and Compact. The default value is Delete. The Delete policy will remove old segments when the retention time or size limit has reached. The Compact policy will enable log compaction on the topic, which ensures that Kafka will always retain at least the last known value for each message key within the log of data for a single topic partition.
- ♦ **-s --schemaregistryport:** Specify the port on which the schema registry is running. The default value is 8081.
- ♦ **-k --kafkaport:** Specify the port on which Kafka listens. The default value is 9093.
- ♦ **-x --connectport:** Specify the port on which Kafka connect listens. The default value is 8083.

For example: `zman server-role-kafka-update-cluster -c=3`

To remove a broker from the cluster or to re-configure a broker, see [Maintaining the Kafka Cluster](#).

After adding a broker to the Kafka cluster, ensure that you run the Configure action to allocate appropriate memory. For more information, see [Managing Vertica Memory Requirements](#).

Vertica

A summary of the steps to add servers to the Vertica cluster are as follows:

1. Execute the command `zman server-role-vertica-add-server (zman srvas)` to add servers to the cluster. Specify the DNS, GUID or path of the server object (server, server folder or server group) relative to /Devices/Servers.

For example: `zman server-role-vertica-add-server --servers=server2.microfocus.com,server3.microfocus.com`

2. Restart the ZENserver and ZENloader services.
3. Execute the command `zman server-role-vertica-update-cluster (zman srvuc)` to update the K factor.

For example: `zman server-role-vertica-update-cluster -k=1`

Adding Servers

- ♦ Execute the following command to add servers to the existing cluster and specify a comma separated list of servers by specifying the DNS, GUID or path of the server object (server, server folder or server group) relative to /Devices/Servers.

`zman server-role-vertica-add-server (zman srvas)`

For example: `zman server-role-vertica-add-server --servers=server2.microfocus.com,server3.microfocus.com`

IMPORTANT: After adding servers to the Vertica cluster, ensure that you restart the ZENserver and ZENloader services on all the Primary Servers.

Updating the K Factor

- ♦ Execute the following command to update the K factor.

```
zman server-role-vertica-update-cluster (zman srvuc)
```

This command lets you update the K-safety factor in the existing cluster. The K-safety value determines the number of replicas of data that should exist in a cluster. The K safety sets the fault tolerance for the Vertica cluster, where K can be set as 0 or 1. Specify a value based on the number of nodes in the cluster, which is measured as $2k+1$. For example: if the K factor is 1, then you need to have a minimum of 3 nodes. Vertica recommends a K-safety value of 1. For more information on K-safety, see [Glossary](#).

For example: `zman server-role-vertica-update-cluster -k=1`

After adding a broker to the Vertica cluster, ensure that you run the Configure action to allocate appropriate memory. For more information, see [Managing Vertica Memory Requirements](#).

Maintaining ZooKeeper, Vertica and Kafka Clusters

The topics covered in this section are:

- ♦ [“Maintaining the ZooKeeper Cluster” on page 29](#)
- ♦ [“Maintaining the Kafka Cluster” on page 30](#)
- ♦ [“Maintaining the Vertica Cluster” on page 32](#)

Maintaining the ZooKeeper Cluster

- ♦ **Modifying the ZooKeeper Cluster:** Execute the following command to modify the cluster level parameters of the ZooKeeper service.

```
zman server-role-zookeeper-update-cluster (zman srzuc)
```

The parameters of this command are:

- ♦ **-l --leader-connect-port:** Define the leader port that the followers will use to connect to a leader in the cluster. The default port is 6790. However, you can specify an unused port between 6000 and 65535.
- ♦ **-c --client-port:** Define the port on which ZooKeeper will listen in for incoming client connections. The default port is 6789. However, you can specify an unused port between 6000 and 65535.
- ♦ **-e --leader-elect-port:** : Define the election port that all servers in the cluster will use to elect the leader. The default port is 6791. However, you can specify an unused port between 6000 and 65535.

- ♦ **-t --tick-time:** Define the length of a single tick, which is a basic time unit in milliseconds, used by ZooKeeper to regulate heartbeats and timeouts. The default value is 2000.
- ♦ **-i --init-limit:** Define the amount of time, in ticks, to allow followers to connect to the leader. The default value is 10.
- ♦ **-s --sync-limit:** Define the amount of time, in ticks, to allow followers to sync with ZooKeeper. The default value is 5.
- ♦ **-x --max-connections:** Specify the maximum number of client connections that the ZooKeeper cluster can accept. The default value is 120.

For example: `zman server-role-zookeeper-update-cluster -s=6`

- ♦ **Removing additional ZooKeeper servers:** Execute the following command to remove additional ZooKeeper servers:

```
zman server-role-zookeeper-remove-server (zman srzrs)
```

This command removes any additional servers from the existing cluster. Specify the DNS, GUID or path of the server object (server, server folder or server group) relative to /Devices/Servers. At least one server in the zone must have a ZooKeeper role enabled on it. Therefore, this command will not work if the server in which ZooKeeper role is enabled, is the last ZooKeeper server.

For example: `zman server-role-zookeeper-remove-server --servers=server1.microfocus.com`

- ♦ **Viewing list of ZooKeeper servers:** Execute the following command to view the list of ZooKeeper servers in a cluster.

```
zman server-role-zookeeper-list-cluster (zman srzlc)
```

Maintaining the Kafka Cluster

Maintaining Kafka Brokers

- ♦ **Removing brokers from the Kafka cluster:** Execute the following command to remove a broker from the Kafka cluster:

```
zman server-role-kafka-remove-broker (zman srkrb)
```

This command lets you remove Kafka brokers from an existing cluster. Specify the DNS, GUID or path of the server object (server, server folder or server group) relative to /Devices/Servers.

For example: `zman server-role-kafka-remove-broker --servers=server1.microfocus.com`

- ♦ **Re-configuring Kafka broker:** Execute the following command to re-configure the Kafka broker:

```
zman server-role-kafka-reconfig-broker (zman srkrb)
```

Specify the DNS, GUID or path of the server object (server, server folder or server group) relative to /Devices/Servers, while executing this command. When a broker is added to Kafka cluster, the remaining servers are re-configured by ZENworks. However, if re-configuration fails for a specific server, then run this command for that server. This command is also used for Disaster Recovery scenarios. For more information, see [ZENworks Disaster Recovery Reference](#)

For example: `zman server-role-kafka-reconfig-broker --servers=server1.microfocus.com`

- ♦ **Viewing list of Kafka brokers:** Execute the following command to view the list of Kafka brokers.

```
zman server-role-kafka-list-cluster (zman srklc)
```

Maintaining Kafka Connect

When Kafka is enabled, Kafka connectors are created that is associated with each table in the RDBMS database. If any of these connectors are not running, then you might have to restart these connectors or re-configure the connectors.

To view the status of these connectors, navigate to the Diagnostic page in ZCC and view the status in the Data Sync Status section.

- ♦ **Restart Kafka Connectors:** Execute the following command:

```
zman server-role-kafka-restart-connectors (zman srkrcn)
```

Specify a comma separated list of connectors. This command can be used when any of the Kafka connectors displayed in the Diagnostics page in ZCC, is not running.

For example: `zman server-role-kafka-restart-connectors -c=zenconnectorzBundlefolderrights,zenconnectorzbundlerights`

- ♦ **Reconfigure Kafka Connectors:** Execute the following command to reconfigure connectors based on the properties mentioned in the connector-configs.xml. This file is available at `etc/opt/novell/zenworks`.

```
zman server-role-kafka-reconfigure-connectors (zman srkccn)
```

The connector-configs.xml file lets you modify properties such as the interval at which the Kafka connector identifies changes in the RDBMS tables or the interval at which updated data is published to Kafka. Specify a comma separated list of connectors.

For example: `zman server-role-kafka-reconfigure-connectors -c=zenconnectorzBundlefolderrights,zenconnectorzbundlerights`

- ♦ **Re-create Kafka Connectors:** Execute the following command to re-create connectors.

```
zman server-role-kafka-recreate-connectors (zman srkrcc)
```

You can execute this command if you have migrated your primary database to another RDBMS database.

For example: `zman server-role-kafka-recreate-connectors -f`

- ♦ **Retrieve Kafka configuration details:** Execute the following command to retrieve the Kafka configuration details:

```
zman server-role-kafka-get-connector-config (zman srkgcc)
```

For example: `zman server-role-kafka-get-connector-config -c=zenconnectorzBundlefolderrights`

- ♦ **Retrieve list of Kafka connectors:** Execute the following command to retrieve a list of Kafka connectors:

```
zman server-role-kafka-list-connectors (zman srklcn)
```

Maintaining the Vertica Cluster

IMPORTANT: If you modify the Vertica cluster, ensure that you restart the ZENserver and ZENloader services on all the Primary Servers.

- ♦ **Removing a server:** Execute the following command to remove a server from the existing cluster:

```
zman server-role-vertica-remove server (zman srvrs)
```

Specify the DNS, GUID or path of the server object (server, server folder or server group) relative to /Devices/Servers, while executing this command. You can only remove one server at a time. Ensure that the K-safety factor remains effective after removing a server from the cluster. If you are decommissioning a server from the existing cluster, then before removing the server, it is recommended that you add another server to the cluster.

For example: `zman server-role-vertica-remove server --servers=server1.microfocus.com`

- ♦ **Preparing the Vertica server for restore process:** Execute the following command to prepare the Vertica server during the backup and restore process. For more information, see [Replacing an Existing Primary Server with a New Primary Server](#).

```
zman server-role-vertica-prepare-server (zman srvps)
```

- ♦ **Viewing list of Vertica servers:** Execute the following command to view the list of Vertica servers in a cluster:

```
zman server-role-vertica-list-cluster (zman srvlc)
```

Monitoring the Status of Vertica, Kafka and ZooKeeper




After configuring Vertica and its related components, you can view the overall configuration status of these components in the [Performance Upgrade](#) page in ZCC. For more information, see [Viewing the Configuration Status](#).

However, if you want to continuously monitor the status of each of the components within ZooKeeper, Kafka and Vertica, see [Monitoring Diagnostics](#).

Viewing the Configuration Status

To view the configuration status of the ZooKeeper, Kafka, Vertica clusters and the Bulk Data Migration process, you can navigate to the Getting Started page in ZCC. To navigate to this page, click [Configuration > Performance Upgrade](#).

Each task on this page, includes an icon that indicates the following status:

- ♦ : the component is successfully configured.
- ♦ : the component is ready to be configured. For example, the Bulk Data Migration component will be ready to be configured only when Kafka and Vertica components are configured.
- ♦ : an error has occurred while configuring the component.

Monitoring Diagnostics

To monitor the overall health of the Vertica, ZooKeeper and Kafka clusters, you need to navigate to the Diagnostics page in ZCC. You can also view the status of the Data Sync process on this page. To navigate to the diagnostics page in ZCC, click **Diagnostics** in the left pane of ZCC.

For more information on debugging issues that are identified by the Diagnostics feature in ZCC, see [Debugging and Log Location](#).

NOTE: If client authentication is not enabled for external CA certificates, then the Data Sync Status and Kafka Cluster status will display correct values only if the Diagnostics page is accessed from a server in which the Kafka role is enabled.

The topics addressed in this section are:

- ♦ [“Vertica Cluster” on page 33](#)
- ♦ [“Zookeeper Cluster” on page 33](#)
- ♦ [“Kafka Cluster” on page 34](#)
- ♦ [“Data Sync Status” on page 34](#)

Vertica Cluster

The statuses displayed are as follows:



Up: indicates that all servers are up.



Down: if at least one server is down in the zone, then based on the K factor the relevant status is displayed. For example, the status is displayed as down if at least one server is down and K factor is 0.



Warning: if at least one server is down in the zone, then based on the K factor the relevant status is displayed. For example, the status is displayed as warning if at least one server is down and K factor is 1.

For debugging issues related to the Vertica cluster, see [Vertica](#)

Zookeeper Cluster

The statuses displayed are as follows:



Up: indicates that majority of servers in the cluster are up.



Down: indicates that majority of servers in the cluster are down.



Warning: if one of the servers is down.

NOTE: Besides Vertica, other features in ZENworks require the ZooKeeper services to be up and running at all times. Therefore, even if you do not have Vertica configured in your zone, ZooKeeper should always be active

For debugging issues related to the ZooKeeper cluster, see [ZooKeeper](#)

Kafka Cluster

The statuses displayed are as follows:



Up: indicates that all servers are up.



Down: if at least one server is down in the zone, then based on the replication count the relevant status is displayed. For example, the status is displayed as down if at least one server is down in a three node cluster with replication count being 1.



Warning: if at least one server is down in the zone, then based on the replication factor the relevant status is displayed. For example, the status is displayed as warning if at least one server is down in a three node cluster with replication factor being 2.

For the Kafka Cluster, the status of the Kafka Brokers, Kafka Connect and the Schema registry are also displayed.

For debugging issues related to the Kafka cluster, see [Debugging and Log Location](#)

Data Sync Status

The Data Sync Status panel enables you to identify whether the data sync process is running seamlessly or not. This panel displays the status of each Kafka connector that is responsible for syncing data from the RDBMS to Vertica. These connectors help in identifying the rows in the database tables that have been updated and publishes this data to Kafka. The ZENworks loader stream processor services then consumes this data from Kafka and sends it to Vertica.

Along with the connector name and its status, this panel displays the name of the ZENloader stream processors (consumer name) and its status and the names of the server on which the loader services are running. It displays the time at which the ZENloader stream processors last consumed data from Kafka and the time at which this data was last sent to Vertica. It also displays the number of pending records that are yet to be migrated from each Kafka connector to Vertica.

For debugging issues related to the Data Sync process, see [Data Migration](#).

Connecting to Vertica

The dbadmin or the zenvertica user can connect to Vertica using the vsql tool. For more information on the vsql tool and to connect to the Vertica database, see the [Vertica](#) documentation.

To connect to Vertica, the user needs to specify the database credentials that can be obtained in either of the following ways:

- ♦ Execute the following zman command:

```
zman server-role-vertica-get-credentials (zman srvgc)
```

- ♦ Obtain the credentials from the `zenworks-coreadmin` web service by referring to the following steps:
 1. On your browser, type the server address of the Primary Server that is part of the Vertica cluster, along with `zenworks-coreadmin`. For example: `https://10.0.0.0/zenworks-coreadmin`.
 2. Select **Test Service**.
 3. Select **Vertica Configuration** and click **Invoke**.

Stopping and Starting Vertica Services

To stop and start the Vertica services, you need to login as a `dbadmin` user by running the following command: `su dbadmin`

- ♦ Execute the following to stop the Vertica services:

```
$ /opt/vertica/bin/admintools -t stop_db -d <db_name> -p <password> -F
```

Use the option `-F` (or `--force`) to override all user connections and force a shutdown:

- ♦ Execute the following command, to start the Vertica services:

```
$ /opt/vertica/bin/admintools -t start_db -d <db_name> -p <password>
```

Specify the `dbadmin` superuser password in both the commands that can be obtained by executing the following:

```
zman server-role-vertica-get-credentials (zman srvgc)
```

To obtain the database name, navigate to the `zenworks-coreadmin` web service by referring to the following steps:

1. On your browser, type the server address of the Primary Server that is part of the Vertica cluster, along with `zenworks-coreadmin`. For example: `https://10.0.0.0/zenworks-coreadmin`.
2. Select **Test Service**.
3. Select **Vertica Configuration** and click **Invoke**. The database name is displayed within the **Database** attribute.

```
zen2020:- # zman srvgc

Vertica zenworks database user details.
Username: zenvertica
Password: Zw0#4f.....

Vertica database superuser details.
Username: dbadmin
Password: Zw0#d70.....

Vertica system user details.
Username: dbadmin
Password: 098L7B1d22P1B11

zen2020:- # su - dbadmin
dbadmin@zen2020:~> /opt/vertica/bin/admintools -t stop_db -d zenworks -p Zw0#d70..... -F
Connecting to database
Issuing shutdown command to database
Database zenworks stopped successfully
dbadmin@zen2020:~> █
```

Shutting Down a Vertica Node

If you need to shut down a Vertica node for maintenance, then ensure that you refer to the following steps to ensure that there is no loss in data:

1. Login to the Vertica database using the dbadmin user and run the following commands:
 - a. `select close_all_sessions();` For more information, see the [Vertica Documentation](#).
 - b. `select make_ahm_now();`
 - c. `select do_tm_task('moveout');`
 - d. `select do_tm_task('mergeout');`
 - e. `select make_ahm_now();`
2. Stop the node using admintools command: `admintools -t stop_node -s <HOSTS>` where `-s <HOSTS>` is a comma separated list of hosts on which the Vertica process needs to be stopped.

6 Disaster Preparedness and Recovery

ZENworks does not recommend maintaining a replica of the Vertica database to protect data in case of a failure in the zone. In the event of a failure of the Primary Server in which Vertica, Kafka or ZooKeeper is enabled, you need to replace the device hosting the Primary Server with a new device that has the same identity, that is the same host name and IP address as that of the old device. For more information, see [Replacing an Existing Primary Server with a New Primary Server](#) in the [ZENworks Disaster Recovery Reference](#).

NOTE: If the replication count is 1, then before decommissioning the last server that has a Kafka role enabled in it, run the `zman server-role-kafka-add-broker (zman srkab)` to add a new broker and then proceed with the decommissioning of the existing server.

However, if you have already decommissioned the server, ensure that you bring up a server with the same hostname by backing up and restoring the Kafka data. For more information, see [Replacing an Existing Primary Server with a New Primary Server](#) in the [ZENworks Disaster Recovery Reference](#).

After replacing the existing Primary Server with a new Primary Server, you need to re-execute the configure action to migrate data from the RDBMS to the Vertica database. For more information, see [“Migrating Data to Vertica” on page 23](#). However, data like the patch CVE status data might not be available immediately unless you run patch scans for all the devices in the zone.

NOTE: In the event of a failure, any trending data that is created on Vertica cannot be recovered. However, Vertica provides a way to take a backup of specific tables in the database and enables you to import these onto another database. For more information, see [Vertica Documentation](#). Therefore, as a precautionary measure, you can take periodic backup of these tables. After replacing the existing Primary Server with the new server, you can replicate and restore these tables from the backup location.

Backing up and Restoring the Vertica Database

As a part of maintenance activities, taking regular backups of the Vertica database is recommended. Vertica provides the `vbr` utility that enables you to perform a full backup and restore operation of the Vertica database.

This utility enables you to perform the following operations:

- ♦ Back up the database
- ♦ Back up specific objects like schemas or tables, in a database.
- ♦ Restore a database or individual objects from a backup.
- ♦ Copy a database to another cluster. For example, to promote a test cluster to production (for Enterprise mode only). For more information, see [Vertica Documentation](#).

- ♦ Replicate individual objects such as schemas and tables to another cluster (for Enterprise mode only). For more information, see [Vertica Documentation](#).
- ♦ List available backups.

The topics addressed in this section are:

- ♦ [“Creating Configuration File for Vertica Full Local Backup” on page 38](#)
- ♦ [“Creating a Full Local Backup” on page 38](#)
- ♦ [“Restoring from a backup” on page 39](#)

Creating Configuration File for Vertica Full Local Backup

To run the vbr utility, you need to specify a configuration file that contains information such as the backup location, name of the database and dbuser, the snapshot name, db administrator password file path, and retry count. Vertica provides a few sample configuration files available at `/opt/vertica/share/vbr/example_configs/backup_restore_full_local.ini` that can be used for various vbr tasks.

If the `dbPromptForPassword` parameter in the configuration file is true, then the vbr utility will prompt for a password every time you run the utility. However, to avoid user intervention while performing the backup operation, you can specify the location of the db administrator password config file in the `passwordFile` parameter. For example, `passwordFile=/path/to/vbr/pw.txt`. This file should be read protected by the db administrator user only. The format of the content in the pw.txt file should be:

[Passwords]

dbPassword=*****

For more information on creating the configuration file, see the [Vertica Documentation](#).

Creating a Full Local Backup

Prerequisites

- ♦ Ensure that you initialize the backup location by running the following command `vbr -t init -c full-backup.ini`.
- ♦ Your database should be up and running. All nodes need not be in a K-safe database. However, the nodes that are down are not backed up.
- ♦ All of the backup hosts should be up and running.
- ♦ Enough disk space should be available to store the backups.
- ♦ The user account of the user who starts the vbr utility should have write access to the target directories on the host backup location. This user can be dbadmin or another assigned role. However, you cannot run the vbr utility as root.
- ♦ Each backup has a unique file name. If you want to keep earlier backups, ensure that the `restorePointLimit` parameter is set to a number greater than 1 in the configuration file

Procedure

Run the vbr utility from a terminal. Use the database administrator account from an initiator node in the database cluster.

Run the following command and a backup will be created without any further intervention:

```
vbr -t backup -c full-backup.ini
```

Restoring from a backup

Prerequisites

Before performing a full restore from the backed up database, ensure that the following requirements are met:

- ♦ The database is down. You cannot perform the restore action when the database is running.
- ♦ All the backup hosts are available.
- ♦ The backup directory exists and contains the backups from which to restore the data.
- ♦ The cluster to which you are restoring the backup should have:
 - ♦ The same number of nodes as the one used to create the backup.
 - ♦ The same architecture (Enterprise or Eon mode) as the one used to create the backup.
 - ♦ Identical node names.

Procedure

To restore from the most recent backup, run the following command:

```
vbr -t restore -c full-backup.ini
```

To restore an archived backup, add the `--archive` parameter to the command line along with the directory name, which is the `date_timestamp` that identifies the archive to restore. For example, `vbr -t restore -c full-backup.ini --archive=20190712_085620`. This will restore from the archive created on 12th July 2019 and 8:56:20.

Changing the ZooKeeper role from the first Primary Server to the second Primary Server

Perform the following steps only if the first Primary Server is up and running. In case of a failure of the first Primary Server, see [Replacing an Existing Primary Server with a New Primary Server](#) in the [ZENworks Disaster Recovery Reference](#).

- 1. Add the second Primary Server to the ZooKeeper cluster:** Run the command `zman srzas --servers=<GUID, path of the server object relative to /Devices/Servers, or the DNS of the second Primary Server>` in the command line utility.

- 2. Remove the first Primary Server:** Run the command `zman srzrs --servers= <GUID, path of the server object relative to /Devices/Servers, or the DNS of the first Primary Server>` in the command line utility.

7 Debugging and Log Location

The topics addressed in this section are:

- ♦ “ZooKeeper” on page 41
- ♦ “Kafka” on page 42
- ♦ “Vertica” on page 43
- ♦ “Data Migration” on page 44
- ♦ “Firewall Issues” on page 44

NOTE: Although ZMAN commands to enable Vertica, Kafka and ZooKeeper can be run on any server, to debug, you need to check the logs on only those servers for which the ZMAN command is executed. For example: if the process to add a server `server1.microfocus.com` to the Kafka cluster has failed, then check the logs on `server1.microfocus.com`.

ZooKeeper

Configuration Issues

For any configuration related issue you can first refer to either one of the logs:

- ♦ `/var/opt/novell/log/zenworks/loader-messages.log`
- ♦ `/var/opt/novell/log/zenworks/zman.log`
- ♦ `%ZENWORKS_HOME%\logs\loader-messages.log`
- ♦ `%ZENWORKS_HOME%\logs\zman.log`

However, if information in these logs is insufficient, then it is recommended that you enable the ZooKeeper logs by uncommenting the lines below the comment `#zookeeper` in the `log4j2.properties` file present in `/opt/novell/zenworks/share/tomcat/conf` and then restarting the ZENworks server. You can then view the `zookeeper-messages.log` in the path `/var/opt/novell/log/zenworks/`.

General Debugging

If at any point in time, the ZooKeeper cluster is down, then check whether the port 6789 of the ZooKeeper server is not being used by any other component.

Kafka

Configuration Issues

For any issue related to adding a Kafka broker, you can refer to either one of the logs on the server on which the Kafka role is to be enabled:

- ♦ `/var/opt/novell/log/zenworks/loader-messages.log`
- ♦ `/var/opt/novell/log/zenworks/zman.log`

General Debugging

Kafka Broker Issues: If the status of **Kafka broker** in the Diagnostics page is **Not Running**, then check whether the port 9093 of the server is not being used by any other component. You can also check the Kafka logs within the `server.logs` file available at `/var/opt/novell/log/zenworks/kafka-logs`.

Schema Registry Issues: If the status of **Schema Registry** in the Diagnostics page is **Not Running**, then check `schema-registry.log` available at `/var/opt/novell/log/zenworks/kafka-logs/var/opt/microfocus/log/zenworks/kafka-logs`.

Kafka Connect Issues: If the status of **Kafka Connect** in the Diagnostics page is **Not Running**, then check whether the port 8083 of the server is not being used by any other component. You can also check the `kafka-connect.log` available at `/var/opt/novell/log/zenworks/kafka-logs`.

If due to some reason any of the Kafka connectors displayed in the **Data Sync** section of the **Diagnostics** page, have failed, then you need to reconfigure the Kafka connector and restart the Kafka-connect service. To reconfigure the connector run the following command:

```
zman server-role-kafka-reconfigure-connectors -c <name of the connector>
```

You can also retrieve the names of all the connectors in the command line utility by running the `zman server-role-kafka-list-connectors` command. After reconfiguring the Kafka connector, restart the Kafka-connect service by running the following command:

```
systemctl restart zenworks-connect.service
```

Also, if the RDBMS is down for more than an hour, the Kafka connectors will not be able to sync data between Vertica and the RDBMS. You need to restart the Kafka.connect service by running the command:

```
systemctl restart zenworks-connect.service
```

```
systemctl restart microfocus-connect.service
```

Vertica

Configuration Issues

For failures related to Vertica cluster configuration or addition of servers to the Vertica cluster, you can check the following logs on the server on which Vertica role is enabled or to be enabled:

Installation Failure

Refer to the following logs:

- ♦ Check `vertica-install.log` within `/var/opt/novell/log/zenworks`
- ♦ Check `vertica-admin-tools.log` within `/var/opt/novell/log/zenworks`
- ♦ Check `zman.log` within `/var/opt/novell/log/zenworks/`
- ♦ Check `vertica-install-prepare.log` within `/var/opt/novell/log/zenworks`
- ♦ Check `vertica-post-database-creation.log` within `/var/opt/novell/log/zenworks`

You can also refer to the following logs for detailed reports:

- ♦ `install.log` available at `/opt/vertica/log`
- ♦ `adminTools.log` available at `/opt/vertica/log`

Vertica Management Console

As the Diagnostics page in ZCC provides only an overview of the status of the Vertica database, to get an in-depth view of the Vertica database operations, you can use the Vertica management console, which is a third-party tool. This tool can be installed and accessed from any server.

For more information on installing the Vertica management console, see the [Vertica Documentation](#).

After logging in to the Management console, you need to select the **Import an existing database/cluster** option. Using this option, you can import multiple Vertica database servers. While importing the database, you need to provide the database credentials, which can be obtained in one of the following ways:

- ♦ Execute the following `zman` command

```
zman server-role-vertica-get-credentials (zman srvgc)
```
- ♦ Obtain the credentials from the `zenworks-coreadmin` web service by referring to the following steps:
 1. On your browser, type the server address of the Primary Server that is part of the Vertica cluster, along with `zenworks-coreadmin`. For example: `https://10.0.0.0/zenworks-coreadmin`.
 2. Select **Test Service**.
 3. Select **Vertica Configuration** and click **Invoke**.

For more information on using the Vertica management console, see the [Vertica Documentation](#).

Data Migration

The logs for bulk data migration is present in the following locations:

- ♦ `novell-zenworks-configure.log` available at `/etc/opt/novell/zenworks`
- ♦ `vertica_migration_csv` available at `/var/opt/novell/zenworks/`. This file provides a more detailed report.

Firewall Issues

If due to firewall restrictions, the other Primary Servers in your zone are unable to access the Vertica, Kafka, or ZooKeeper services, then you need to run the following Configure action on the server in which these components are enabled, to open the ports. However, if the server in the DMZ is unable to access the ZooKeeper services within the corporate network, then you need to manually open the port 6789 in the corporate firewall.

Example:

If you are unable to access the ZooKeeper services, to open the ZooKeeper ports, run the following configure action:

```
novell-zenworks-configure -c ClusterFirewallConfigureAction -  
Doperation=add -Dservice=zookeeper
```

If you are unable to access the Kafka services, to open the Kafka ports, run the following configure action:

```
novell-zenworks-configure -c ClusterFirewallConfigureAction -  
Doperation=add -Dservice=kafka
```

If you are unable to access the Vertica services, to open the Vertica ports, run the following configure action:

```
novell-zenworks-configure -c ClusterFirewallConfigureAction -  
Doperation=add -Dservice=vertica
```

To close the ports, you need to run the following Configure action:

Example:

To close the ZooKeeper ports, run the following configure action:

```
novell-zenworks-configure -c ClusterFirewallConfigureAction -  
Doperation=remove -Dservice=zookeeper
```

To close the Kafka ports, run the following configure action:

```
novell-zenworks-configure -c ClusterFirewallConfigureAction -  
Doperation=remove -Dservice=kafka
```

To close the Vertica ports, run the following configure action:

```
novell-zenworks-configure -c ClusterFirewallConfigureAction -  
Doperation=remove -Dservice=vertica
```

For more information on the default ports that need to be open to access the ZooKeeper, Kafka and Vertica services, see [ZENworks 2020 TCP and UDP Ports](#).

8 Removing Vertica

This section details the steps to disable Vertica in the zone. Prior to disabling Vertica, it is important that you are aware of the following implications of performing this procedure:

- ♦ Vertica will be completely removed from the zone
- ♦ You will not have the option to revert to the Vertica database
- ♦ All information related to the Trending dashlets will be removed

If you want to re-configure Vertica in the zone, then you need to start all over again by enabling Vertica as detailed in the section [Configuring Vertica](#).

- ♦ [“Disabling and Removing Vertica from the Zone” on page 47](#)

Disabling and Removing Vertica from the Zone

To disable Vertica:

- 1 Run the following zman command on any server to obtain the Vertica database superuser password.

```
server-role-vertica-get-credentials (zman srvgc)
```

- 2 Using the obtained dbadmin password, run the following command to stop the Vertica database on one of the servers on which Vertica is installed:

```
/opt/vertica/bin/admintools -t stop_db -d zenworks -p  
<dbadmin_password> --force.
```

- 3 Run the following drop database command on one of the servers on which Vertica is installed:

```
/opt/vertica/bin/admintools -t drop_db -d zenworks
```

- 4 Stop all the cluster nodes. To identify these nodes, run the following command on one of the servers on which Vertica is installed:

```
/opt/vertica/bin/admintools -t list_host
```

If a node is listed, then run the following command to stop the node:

```
/opt/vertica/bin/admintools -t stop_node -s <comma separated list of  
hosts>.
```

- 5 Disable the following services on the respective Vertica and Kafka servers.

zenworks-connect.service: Execute the following commands in the order listed below:

```
systemctl stop zenworks-connect.service
```

```
systemctl disable zenworks-connect.service
```

zenworks-kafka.service: Execute the following commands in the order listed below:

```
systemctl stop zenworks-kafka.service
```

```
systemctl disable zenworks-kafka.service
```

zenworks-schema-registry.service: Execute the following commands in the order listed below:

```
systemctl stop zenworks-schema-registry.service
```

```
systemctl disable zenworks-schema-registry.service
```

zenworks-vertica.service: Execute the following commands in the order listed below:

```
systemctl stop zenworks-vertica.service
```

```
systemctl disable zenworks-vertica.service
```

- 6 Delete the following datamodel files on all the servers.

On a Linux server: Navigate to `/etc/opt/novell/zenworks/datamodel` and delete the files `zenvertica_dmaccounts.properties` and `zenvertica.xml`. Also, remove the entry `vertica=zenvertica` from the `dmmappings.properties` file.

On a Windows server: Navigate to `C:\Program Files (x86)\Novell\ZENworks\conf\datamodel` and delete the files `zenvertica_dmaccounts.properties` and `zenvertica.xml`. Also, remove the entry `vertica=zenvertica` from the `dmmappings.properties` file.

- 7 Run the following query to remove the opaque entries related to Vertica from the database:

```
DELETE FROM ZOPAQUEDATA WHERE name LIKE '%vertica%'
```

- 8 Run the following query to remove the opaque entries related to Kafka from the database:

```
DELETE FROM ZOPAQUEDATA WHERE name LIKE '%kafka%'
```

- 9 Remove the following cluster ID file from all the Vertica nodes to stop the stream processors:

```
/etc/opt/novell/zenworks/loader/loader.cluster.id.
```

- 10 Remove the Vertica role from the database:

```
DELETE FROM ZZENSERVERRoles roles WHERE Roles = 'Vertica'
```

- 11 Remove the Kafka role from the database:

```
DELETE FROM ZZENSERVERRoles roles WHERE Roles = 'Kafka'
```

- 12 Remove the Kafka connector information from the database:

```
DELETE FROM ZVCONSUMERMETRICS;
```

- 13 Restart the ZENserver and ZENloader services on all the servers.

Enabling Vertica

1. Procure the obfuscated license file from Micro Focus Customer Support and copy the file in the location `/etc/opt/novell/vertica`.
2. Enable Kafka.
3. Enable Vertica. For more information, see [Enabling Vertica](#).

4. Execute the data migration configure action using the Force option:

```
novell-zenworks-configure -c VerticaDBMigrate -Doption=force
```


9 Troubleshooting

This section provides solutions to problems that you might face with the Vertica, Kafka or ZooKeeper components.

ZooKeeper

An Apache curator error message is displayed in the logs, when new servers are added to the ZooKeeper cluster

Explanation: When additional servers are added to the ZooKeeper cluster, the following Apache curator related error message is displayed in the logs for each newly added server, even though the servers are successfully added in the cluster:

ERROR org.apache.curator.framework.imps.EnsembleTracker - Invalid config event received

Action: As there is no loss in functionality and the servers are added successfully, no action needs to be performed. However, a bug has been raised with Apache.

Kafka

Addition of another server to the Kafka cluster fails after the only server that was part of the Kafka cluster is decommissioned

Explanation: Consider a scenario, where the replication count is 1 and the only server (broker) that was added to the Kafka cluster is decommissioned. On executing the command `zman server-role-kafka-add-broker (zman srkab)` to add another server (broker) in the zone, the operation to add the broker to the Kafka cluster fails. The error messages **Schema registry started – false** and **Some of the processes have failed to start. Marking queue as failed** are displayed in the `zenworks loader-messages.log`.

Action: Refer to the following steps to clean up both the ZooKeeper and Kafka data directory:

For ZooKeeper:

1. Stop all the ZENworks services on the Primary Server on which the ZooKeeper role is enabled. For more information, see [Stopping the ZENworks Services](#).
2. Remove the `zookeeper.properties` file from the path `%ZENWORKS_HOME%\novell\ZENworks\conf` (on Windows) or `/etc/opt/novell/ZENworks/conf` (on Linux).

3. Remove the ZooKeeper data directory from the path
`%ZENWORKS_HOME%\novell\ZENworks\work\zookeeper` (on Windows) or `/var/opt/novell/ZENworks/zookeeper` (on Linux).
4. Restart all the ZENworks services. For more information, see [Starting the ZENworks Services](#).

For Kafka:

1. Stop all the ZENworks services on the server on which the `zman` command to add the Kafka broker was executed but failed. For more information, see [Stopping the ZENworks Services](#).
2. Remove the Kafka data directory from the path `/var/opt/novell/zenworks/kafka-data`.
3. Restart all the ZENworks services. For more information, see [Starting the ZENworks Services](#).

NOTE: If the replication count is 1, then before decommissioning the last server that has a Kafka role enabled in it, run the `zman server-role-kafka-add-broker (zman srkab)` to add a new broker and then proceed with the decommissioning of the existing server.

However, if you have already decommissioned the server, ensure that you bring up a server with the same hostname by backing up and restoring the Kafka data. For more information, see [Replacing an Existing Primary Server with a New Primary Server](#) in the [ZENworks Disaster Recovery Reference](#).

At least one Kafka connector has failed

Explanation: If for any reason any of the Kafka connectors have failed, which is displayed in the **Data Sync** panel of the Diagnostics page in ZCC, then the failed connectors will not be able to sync data from their associated RDBMS tables to Vertica.

Action: Run the command `zman server-role-kafka-reconfigure-connectors -c <name of the connector>` to reconfigure the failed Kafka connector.

You can retrieve the names of the all connectors within the command line utility by running the `zman server-role-kafka-list-connectors` command.

After reconfiguring the Kafka connectors, you need to restart the Kafka-connect service by running the following command:

```
systemctl restart zenworks-connect.service
```

```
systemctl restart microfocus-connect.service
```

Vertica

After partially disabling Vertica in the zone, upgrade to the next ZENworks release version fails

Explanation: Consider a scenario, where Vertica is configured on a Primary Server, and then needs to be disabled and removed from the zone. While disabling Vertica, a few errors occurred, due to which Vertica was only partially disabled. Subsequently, upgrade to the next ZENworks release version fails.

Action: Before proceeding with the Upgrade process, ensure that you have fully configured or fully disabled Vertica. While disabling Vertica in your zone, ensure that all the steps documented in [Disabling and Removing Vertica from the Zone](#) are followed and no errors occur while executing the commands.

However, if you are still unable to disable Vertica completely from the zone, then run the following query:

```
delete from preglobalupgradelog where schematype='vertica'
```

A node in the Vertica cluster shuts down

Explanation: To provide fault tolerance, Vertica can be run as a cluster of nodes. However, at times the nodes in the cluster can abruptly shut down due to issues related to network connectivity and so on.

Action: Follow the steps documented in the [Vertica](#) documentation, to bring up the database when one of the nodes shuts down.

When nodes responsible for bringing a Vertica cluster down are restarted, then the remaining nodes are no longer a part of the cluster

Explanation: The K-factor specified while configuring Vertica, sets the fault tolerance for the Vertica cluster, where K can be set as 0 or 1. The Vertica cluster will be down, if more than 'K' number of nodes are down. For example, if the K factor is 1 in a 3 node cluster and 2 nodes are down, then the Vertica cluster will be down. When you restart services on these 2 nodes, the third node (that was up and running) will no longer be a part of the cluster.

Action: The command to restart the Vertica servers should be run for all the servers in the Vertica cluster. For more information on how to restart the servers in the Vertica cluster, see the [Vertica](#) documentation.

A

Appendix 1: Mounting a New Hard Disk for Vertica Catalog

When one or more nodes are running on low disk space or have a failed disk, then beyond a certain threshold, Vertica starts rejecting transactions that update the catalog or data files. In Vertica, catalog is a set of files that contains information (metadata) about the objects in a database, such as the nodes, tables, constraints, and projections. The catalog files are replicated on all nodes in a cluster. To avoid low disk space or failed disk related issues, you can mount a new hard disk and store the catalog files on it.

Refer to the following steps to mount a new hard disk. You can ignore these steps if you have already mounted a disk.

1. Obtain the name of the new hard disk. You can use the command `fdisk -l`.
2. Format the new hard disk. For example, if the name of the new hard disk is `sdc`, then run the command `mkfs.ext4 /dev/sdc`.
3. Create a new directory and mount the hard disk to it. For example, if the new directory is `/vastorage/home/dbadmin/zenworks/NewDisk` then use the commands `mkdir -p /vastorage/home/dbadmin/zenworks/NewDisk` and `mount /dev/sdc /vastorage/home/dbadmin/zenworks/NewDisk`.

Refer to the following steps to move the existing data to the new directory location and make it the primary storage location for Vertica. The location `/vastorage/home/dbadmin/zenworks/NewDisk` has been used as an example in the following steps. Replace this location with the location where the new hard disk has been mounted:

1. Create a new directory within the mounted location and give permissions to it:

```
mkdir -p /vastorage/home/dbadmin/zenworks/NewDisk/  
v_zenworks_node0004_data/ && chown -R dbadmin:verticadba /vastorage/  
home/dbadmin/zenworks/NewDisk/v_zenworks_node0004_data/
```

NOTE: In the above command, `/vastorage/home/dbadmin/zenworks/NewDisk/` is the directory location where the hard disk is mounted and `v_zenworks_node0004_data` is the new directory.

2. Create a new db location by running the following SQL query in Vertica using any connection tool and specify the node to which the data should be moved.

```
CREATE LOCATION '/vastorage/home/dbadmin/zenworks/NewDisk/  
v_zenworks_node0004_data/' NODE 'v_zenworks_node0001' USAGE 'DATA,TEMP'  
LABEL 'TO_DATA_TEMP';
```

NOTE: If you have more than one node, run this query for all the nodes. For example: CREATE LOCATION '/vastorage/home/dbadmin/zenworks/NewDisk/v_zenworks_node0005_data/' NODE 'v_zenworks_node0002' USAGE 'DATA,TEMP' LABEL 'T0_DATA_TEMP';, indicates that directory /vastorage/home/dbadmin/zenworks/NewDisk/v_zenworks_node0005_data/ is mounted in node 2.

3. Set object storage policy at db level by running the following query:

```
select set_object_storage_policy('zenworks','T0_DATA_TEMP', true );
```

4. Apply the storage policy by running the following query:

```
select enforce_object_storage_policy('zenworks');
```

5. Ensure that the new storage location is created successfully, by running the following queries:

To check the storage location:

```
select * from storage_locations
```

To check the storage label:

```
select node_name,location_label from storage_containers where projection_id in (select projection_id from projections);
```

To check whether the new storage location is created or not (the older location will also be displayed till it is retired):

```
select is_retired,location_path from storage_locations;
```

6. Retire the older location by running the following query:

```
SELECT RETIRE_LOCATION('/vastorage/home/dbadmin/zenworks/v_zenworks_node0001_data/' , 'v_zenworks_node0001', true);
```

If you have more than one node, then repeat this for all the older locations of the nodes.

To verify whether the location is retired, check if retire is marked True by running the following query:

```
select is_retired,location_path from storage_locations;
```

The retired location will be displayed as True till the location is deleted.

7. Delete the retired storage location by running the following query:

```
SELECT DROP_LOCATION('/vastorage/home/dbadmin/zenworks/v_zenworks_node0001_data/' , 'v_zenworks_node0001');
```

To check whether the location is deleted, run the following query:

```
select is_retired,location_path from storage_locations;
```

After completing this procedure, it is recommended that you restart the database.

Glossary

Kafka Cluster. A group of servers or nodes that are connected to each other to achieve a common objective is a cluster. Each of the servers or nodes in the cluster, will have one instance of Kafka broker running.

Kafka Brokers. One or more servers that are added in a Kafka cluster are called Brokers. These are typically message brokers that act as mediators between two systems, ensuring that messages are delivered to the correct systems.

Kafka Topics. Kafka stores streams of records in categories called Topics. These topics are further broken down into partitions that can be stored across multiple servers and disks. Topics in Kafka are always multi-subscriber; that is, a topic can have zero, one, or many consumers that subscribe to the data written to it.

Kafka Partitions. Kafka topics are further broken down into partitions that enables topics to be replicated across multiple brokers for fault tolerance.

Kafka Producers. The processes that publish messages to one or more Kafka topics are called Producers. In this case, the producer is the RDBMS system.

Kafka Consumers. The processes that consumes messages from the Kafka topics are called consumers. In this case, the consumer is Vertica.

Kafka Replicas. A replica or a copy of a partition is essentially used to prevent loss of data.

Kafka Leader. A node that handles all read and write requests for a partition is called a leader.

Kafka Follower. A node that replicates the leader are called followers. If the leader fails, one of the followers will automatically become a leader.

Kafka Connect. Kafka Connect is a tool included with Kafka that imports and exports data to Kafka. It is an extensible tool that runs connectors, which implement the custom logic for interacting with an external system. For example, a connector to an RDBMS system captures every change in data made in a table.

Kafka Schema Registry. Kafka producers write data to Kafka topics and Kafka consumers read data from Kafka topics. Schema Registry helps ensure that producers write data with a schema that can be read by consumers, even as producers and consumers evolve their schemas.

Heartbeat Requests. ZooKeeper sends heartbeat requests to determine whether the Kafka broker is alive.

Tick Time. The unit of time used by ZooKeeper that is translated to milliseconds. All time-dependent operations for ZooKeeper are based on tick time.

Offset. A Kafka topic receives messages across a set of partitions and each partition maintains these messages in a sequential order, which is identified as an offset, also known as a position.

K Factor. Specify the number of replicas of the data that should exist in the database. The K safety sets the fault tolerance for the Vertica cluster, where K can be set as 0 or 1. Specify a value based on the number of nodes in the cluster, which is measured as $2k+1$. For example: the K factor 0 is ideal for a single node cluster. If the K factor is 1, then you need to have a minimum of 3 nodes. Vertica recommends a K-safety value of 1.

K-factor also determines the number of identical instance of the segmented projections to be created for fail safety. Vertica stores table data in projections that are a collection of table columns, which optimizes query execution. The projection data can be divided into multiple segments or can be maintained as a single unsegmented unit. To ensure high availability of data, unsegmented projections are replicated on all the nodes and a segmented projection should have $K+1$ identical instances (buddies). For example: if the K- factor is specified as 1, then 2 identical instances of the segment are maintained.

B Documentation Updates

This section contains information on documentation content changes that were made in this *Vertica Reference* for ZENworks Configuration Management. The information can help you to keep current on updates to the documentation.

The documentation for this product is provided on the Web in two formats: HTML and PDF. The HTML and PDF documentation are both kept up-to-date with the changes listed in this section.

If you need to know whether a copy of the PDF documentation that you are using is the most recent, the PDF document includes a publication date on the title page.

The documentation was updated on the following date:

October 2020: ZENworks 2020 Update 1

Location	Change
Chapter 1, "Overview," on page 7	Re-phrased the following statement "You have more than 10 thousand devices in the zone and if you use the Common Vulnerabilities and Exposure (CVE) feature. In such a scenario as significant amount of data is stored in the database it is mandatory to have the Vertica database installed in the zone"
Chapter 8, "Removing Vertica," on page 47	In this section, removed "Disabling Vertica and Reverting to the RDBMS".

