

Planning Your Streams Messaging Deployment

Date published: 2019-12-16

Date modified: 2024-04-03

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with a stylized 'E' that has a horizontal bar extending to the right.

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Data Hub cluster definitions.....	4
Streams Messaging cluster layout.....	4
Streams Messaging: Light Duty cluster layout.....	4
Streams Messaging: Heavy Duty cluster layout.....	7
Streams Messaging: High Availability cluster layout.....	9

Data Hub cluster definitions

The Streams Messaging templates include Kafka, Schema Registry, Streams Messaging Manager, Streams Replication Manager and ZooKeeper. You may choose from the following template options, depending on your operational objectives:

- Streams Messaging Heavy Duty for AWS
- Streams Messaging Light Duty for AWS
- Streams Messaging High Availability for AWS
- Streams Messaging Heavy Duty for Azure
- Streams Messaging Light Duty for Azure
- Streams Messaging High Availability for Azure (Technical Preview)
- Streams Messaging Heavy Duty for GCP
- Streams Messaging Light Duty for GCP
- Streams Messaging High Availability for GCP (Technical Preview)

Streams Messaging provides advanced messaging and real-time processing on streaming data using Apache Kafka, centralized schema management using Schema Registry, management and monitoring capabilities powered by Streams Messaging Manager, as well as cross-cluster Kafka topic replication using Streams Replication Manager and Kafka partition rebalancing with Cruise Control.

These templates set up fault-tolerant standalone deployments of Apache Kafka and supporting Cloudera components (Schema Registry, Streams Messaging Manager, Streams Replication Manager and Cruise Control), which can be used for Kafka workloads in the cloud or as a disaster recovery instance for on-premises Kafka clusters.

Streams Messaging cluster layout

The Data Hub service includes three default Streams Messaging cluster definitions. These are the Streams Messaging: Light Duty, Streams Messaging: Heavy Duty cluster, and Streams Messaging: High Availability definitions. Learn about the layout, capacity, and components of these definitions.

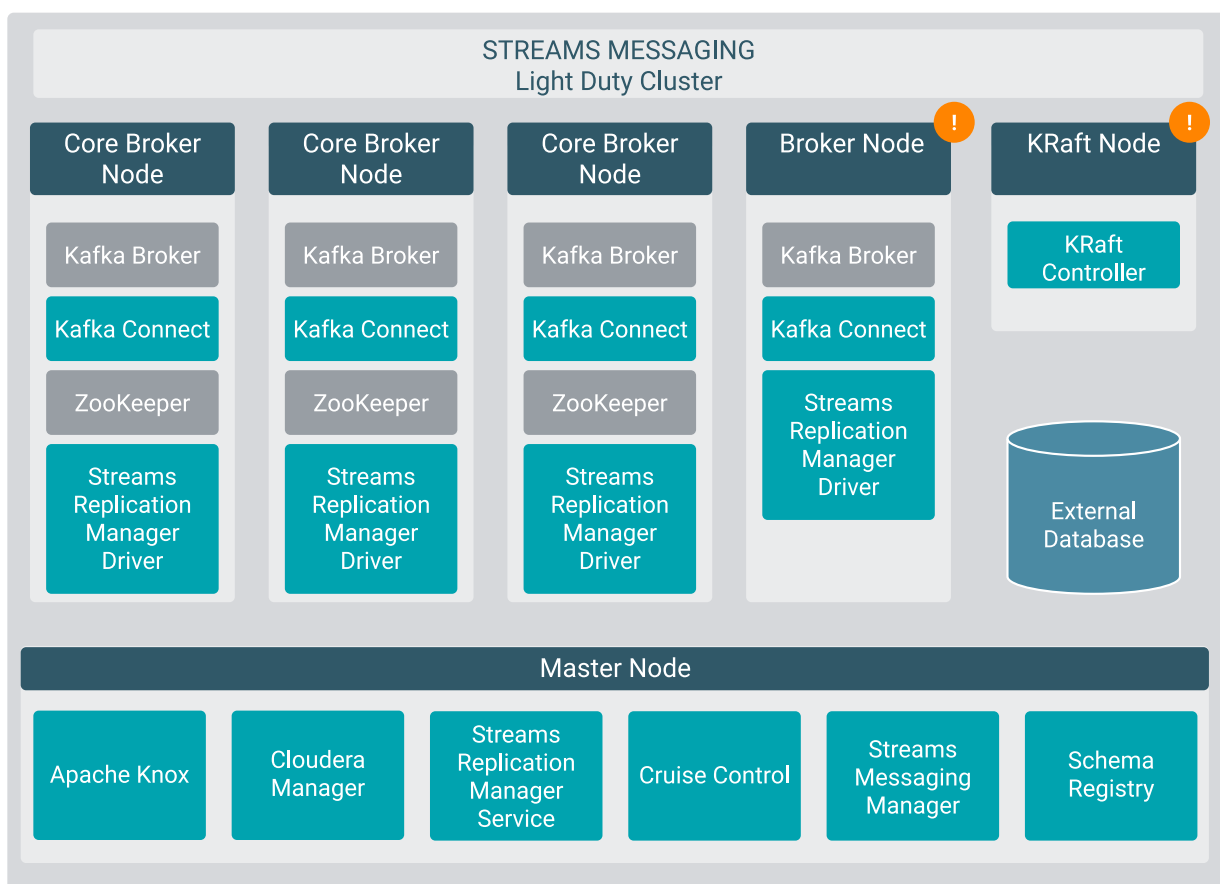


Note: Kafka KRaft is available in this version of CDP but is not ready for production deployment. Cloudera encourages you to explore this technical preview feature in non-production environments and provide feedback on your experiences through the [Cloudera Community Forums](#). For more information regarding KRaft limitations and unsupported features, see [Unsupported Streams Messaging features](#).

Streams Messaging: Light Duty cluster layout

Learn about the layout, capacity, and components of the Streams Messaging: Light Duty cluster definition.

You can use a Streams Messaging: Light Duty cluster definition in development, testing, or proof of concept scenarios. Light Duty clusters include the following nodes and components (services):



Non-default nodes and scaling

Broker and KRaft nodes are not provisioned by default. You have the option to manually set how many of these nodes are created when provisioning the cluster. After the cluster is provisioned, the number of Broker and KRaft nodes can be changed by scaling your cluster. For more information about scaling KRaft and Broker Nodes, see *Scaling Streams Messaging Clusters*.

KRaft and Zookeeper

By default, Kafka uses Zookeeper as its metadata store. If you provision KRaft nodes in the cluster, Kafka runs in KRaft mode and uses KRaft as its metadata store. In this case, ZooKeeper instances are still provisioned. However, ZooKeeper is not used by Kafka to store and manage metadata. If required, ZooKeeper can be removed from the cluster after the provisioning is finished. ZooKeeper can be removed in Cloudera Manager. For more information, see *Deleting ZooKeeper from Streams Messaging clusters*

You can only provision an odd number of KRaft nodes. This is required so that KRaft can hold a majority election for leadership. While it is possible to run Kafka in KRaft mode with a single KRaft node, Cloudera recommends that you provision a minimum of three to avoid having a single point of failure. Cluster provisioning fails if an even number of KRaft nodes are provisioned.

Broker and Core Broker volume per instance count

By default, the volume per instance count for Broker and Core Broker nodes is identical. If you customize your cluster during provisioning, Cloudera recommends that Attached Volume per Instances is set to the same value for both node types. Alternatively, if you want to provision a cluster where the number of volumes is not identical, ensure that you complete *Configure data directories for clusters with custom disk configurations* after the cluster is

provisioned. Otherwise, Kafka does not utilize all available volumes. Additionally, scaling the cluster might not be possible.

Default instance and storage configuration

The following table collects the default instance type and storage configuration of the various nodes deployed with the Streams Messaging: Light Duty cluster. For more information about the cloud provider-specific instance and storage types, see the *Related Information* section.

For Azure

Table 1: Streams Messaging: Light Duty default hardware configuration in Azure

Node	Instance type	Storage configuration
Master	Standard_E8s_v3	100 GB Standard Locally-redundant SSD storage
Core Broker	Standard_D8s_v3	1 TB Locally-redundant storage
Broker	Standard_D8_v3	1 TB Locally-redundant storage
KRaft	Standard_D8_v3	100 GB Locally-redundant storage

For AWS

Table 2: Streams Messaging: Light Duty default hardware configuration in AWS

Node	Instance type	Storage configuration
Master	r5.2xlarge	100 GB Magnetic
Core Broker	m5.2xlarge	1 TB Throughput Optimized HDD
Broker	m5.2xlarge	1 TB Throughput Optimized HDD
KRaft	m5.2xlarge	100 GB Magnetic

For GCP

Table 3: Streams Messaging: Light Duty default hardware configuration in GCP

Node	Instance type	Storage configuration
Master	e2-highmem-8	100 GB Standard Persistent Disk (HDD)
Core Broker	e2-standard-8	1 TB Standard persistent disks (HDD)
Broker	e2-standard-8	1 TB Standard persistent disks (HDD)
KRaft	e2-standard-8	100 GB Standard Persistent Disk (HDD)

Related Information

[Scaling Streams Messaging Clusters](#)

[AWS instance types](#)

[Azure instance types](#)

[GCP instance types](#)

[AWS storage information](#)

[Azure storage information](#)

[GCP storage information](#)

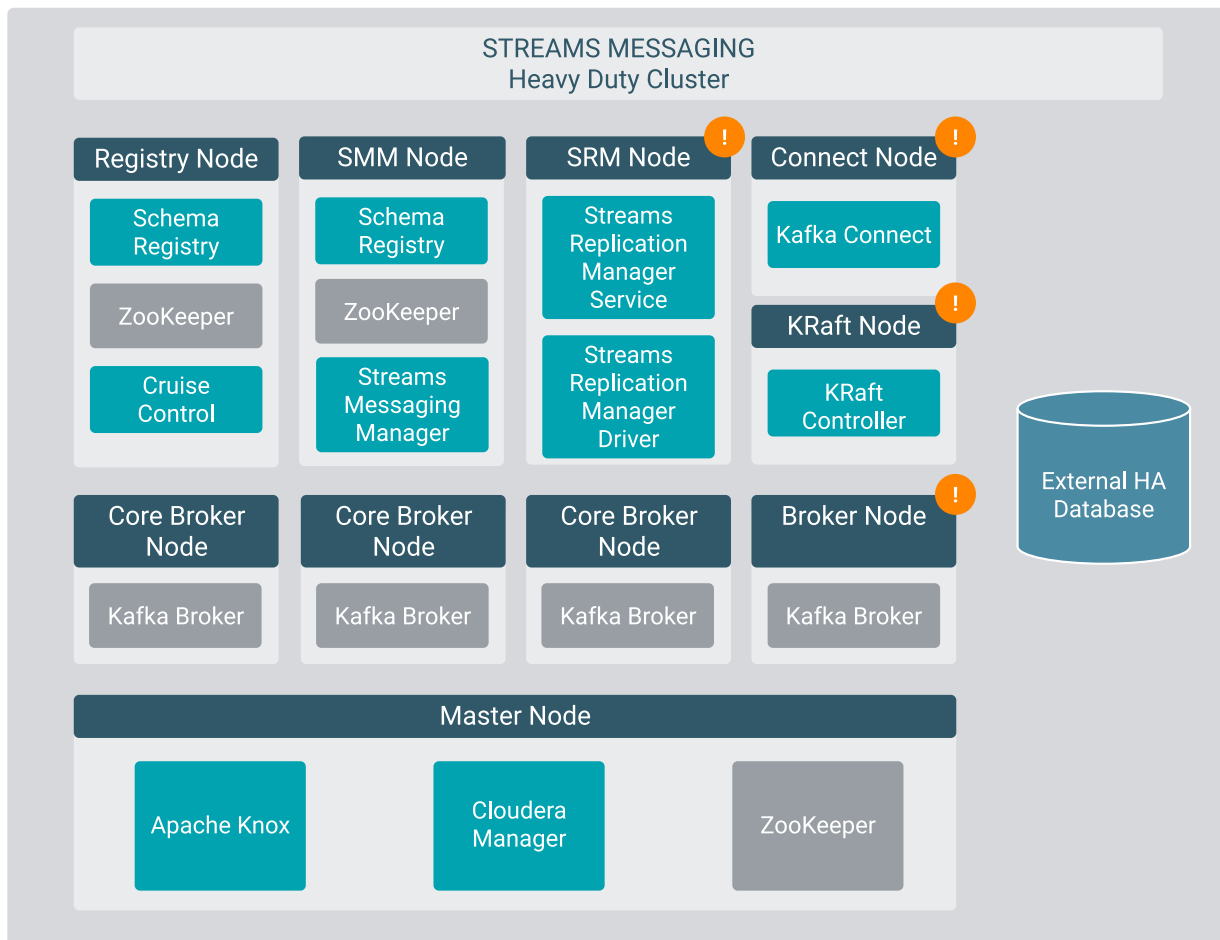
[Configure data directories for clusters with custom disk configurations](#)

[Deleting ZooKeeper from Streams Messaging clusters](#)

Streams Messaging: Heavy Duty cluster layout

Learn about the layout, capacity, and components of the Streams Messaging: Heavy Duty cluster definition.

You can use the Streams Messaging: Heavy Duty cluster definition in production scenarios. Heavy Duty clusters include the following nodes and components (services):



Non-default nodes and scaling

SRM, Broker, Connect, and KRaft nodes are not provisioned by default. If you want to have any of these services provisioned, you must manually set the instance count of the appropriate host group to at least one during cluster provisioning. Otherwise, the host group and its nodes are not provisioned. After a cluster is provisioned, you also have the option to scale these nodes. For more information on scaling, see *Scaling Streams Messaging Clusters*.

KRaft and Zookeeper

By default, Kafka uses Zookeeper as its metadata store. If you provision KRaft nodes in the cluster, Kafka runs in KRaft mode and uses KRaft as its metadata store. In this case, ZooKeeper instances are still provisioned. However, ZooKeeper is not used by Kafka to store and manage metadata. If required, ZooKeeper can be removed from the cluster after the provisioning is finished. ZooKeeper can be removed in Cloudera Manager. For more information, see *Deleting ZooKeeper from Streams Messaging clusters*

You can only provision an odd number of KRaft nodes. This is required so that KRaft can hold a majority election for leadership. While it is possible to run Kafka in KRaft mode with a single KRaft node, Cloudera recommends that you

provision a minimum of three to avoid having a single point of failure. Cluster provisioning fails if an even number of KRaft nodes are provisioned.

Broker and Core Broker volume per instance count

By default, the volume per instance count for Broker and Core Broker nodes is identical. If you customize your cluster during provisioning, Cloudera recommends that Attached Volume per Instances is set to the same value for both node types. Alternatively, if you want to provision a cluster where the number of volumes is not identical, ensure that you complete *Configure data directories for clusters with custom disk configurations* after the cluster is provisioned. Otherwise, Kafka does not utilize all available volumes. Additionally, scaling the cluster might not be possible.

Default instance and storage configuration

The following table collects the default instance type and storage configuration of the various nodes deployed with the Streams Messaging: Heavy Duty cluster. For more information about the cloud provider-specific instance and storage types, see the *Related Information* section.

For Azure

Table 4: Streams Messaging: Heavy Duty default hardware configuration in Azure

Node	Instance type	Storage configuration
Master	Standard_E8s_v3	100 GB Standard Locally-redundant SSD storage
Core Broker	Standard_D8s_v3	1 TB Premium locally-redundant storage
Broker	Standard_D8s_v3	1 TB Premium locally-redundant storage
Registry	Standard_D8_v3	100 GB Locally-redundant storage
SMM	Standard_D8_v3	100 GB Locally-redundant storage
SRM	Standard_D8_v3	100 GB Locally-redundant storage
Connect	Standard_D8_v3	100 GB Locally-redundant storage
KRaft	Standard_D8_v3	100 GB Locally-redundant storage

For AWS

Table 5: Streams Messaging: Heavy Duty default hardware configuration in AWS

Node	Instance type	Storage configuration
Master	r5.2xlarge	100 GB Magnetic
Core Broker	m5.2xlarge	1 TB General Purpose (GP3 SSD)
Broker	m5.2xlarge	1 TB General Purpose (GP3 SSD)
Registry	m5.2xlarge	100 GB Magnetic
SMM	m5.2xlarge	100 GB Magnetic
SRM	m5.2xlarge	100 GB Magnetic
Connect	m5.2xlarge	100 GB Magnetic
KRaft	m5.2xlarge	100 GB Magnetic

For GCP

Table 6: Streams Messaging: Heavy Duty default hardware configuration in GCP

Node	Instance type	Storage configuration
Master	e2-highmem-8	100 GB Standard persistent disks (HDD)
Core Broker	e2-standard-8	1 TB Solid-state persistent disks (SSD)
Broker	e2-standard-8	1 TB Solid-state persistent disks (SSD)
Registry	e2-standard-8	100 GB Standard persistent disks (HDD)
SMM	e2-standard-8	100 GB Standard persistent disks (HDD)
SRM	e2-standard-8	100 GB Standard persistent disks (HDD)
Connect	e2-standard-8	100 GB Standard persistent disks (HDD)
KRaft	e2-standard-8	100 GB Standard persistent disks (HDD)

Related Information

[Scaling Streams Messaging Clusters](#)

[AWS instance types](#)

[Azure instance types](#)

[GCP instance types](#)

[AWS storage information](#)

[Azure storage information](#)

[GCP storage information](#)

[Configure data directories for clusters with custom disk configurations](#)

[Deleting ZooKeeper from Streams Messaging clusters](#)

Streams Messaging: High Availability cluster layout

Learn about the layout, capacity, and components of the Streams Messaging: High Availability cluster definition.

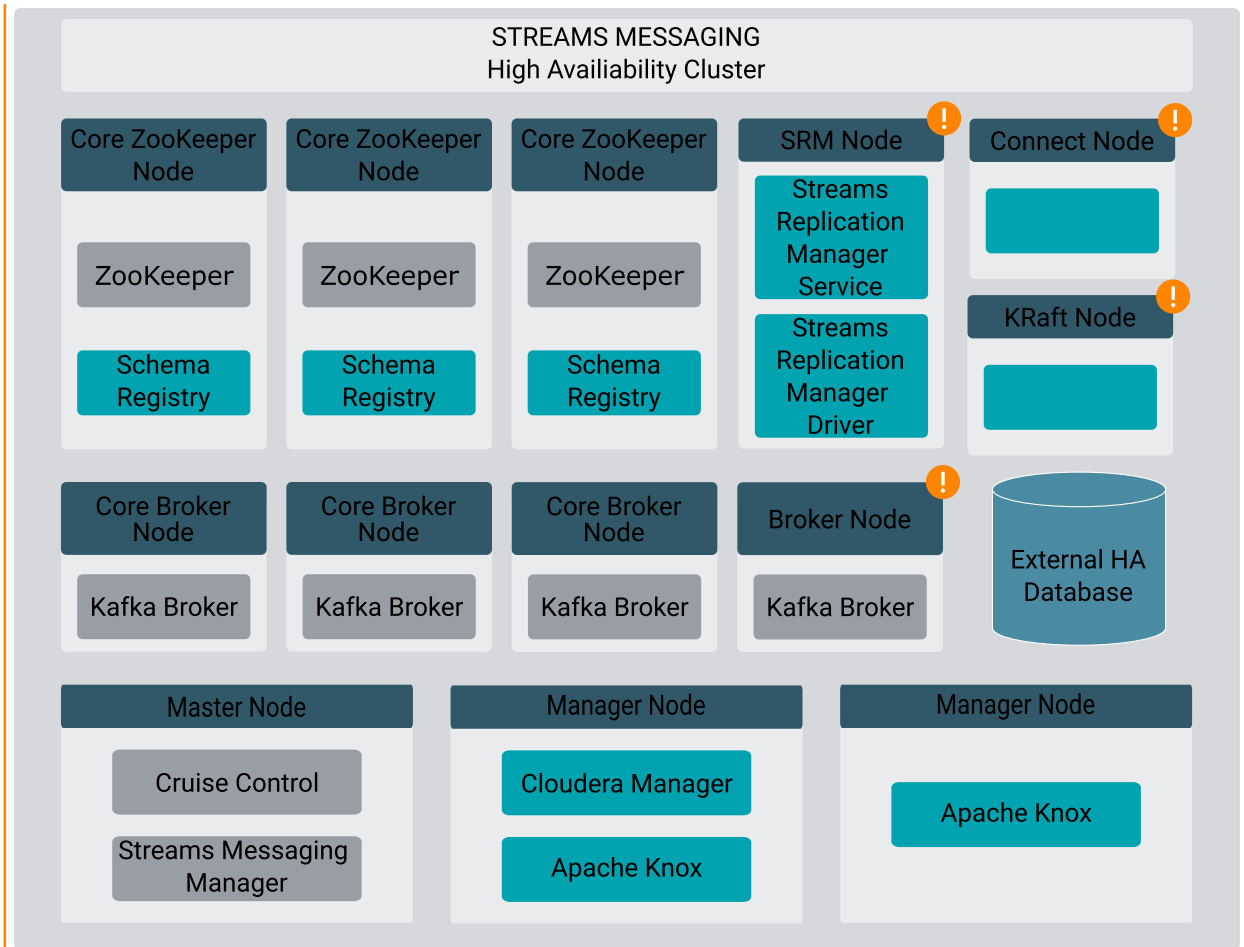


Important: Multi-AZ deployments are not supported for Azure or Google Cloud. As a result, the Streams Messaging High Availability cluster definitions for Azure and Google Cloud are in technical preview. Clusters provisioned with these definitions can only be deployed in a single availability zone.

You can use the Streams Messaging: High Availability cluster definition in production scenarios where having a highly available cluster spanning multiple Availability Zones (multi-AZ) is required. High Availability clusters include the following nodes and components (services):

For V2

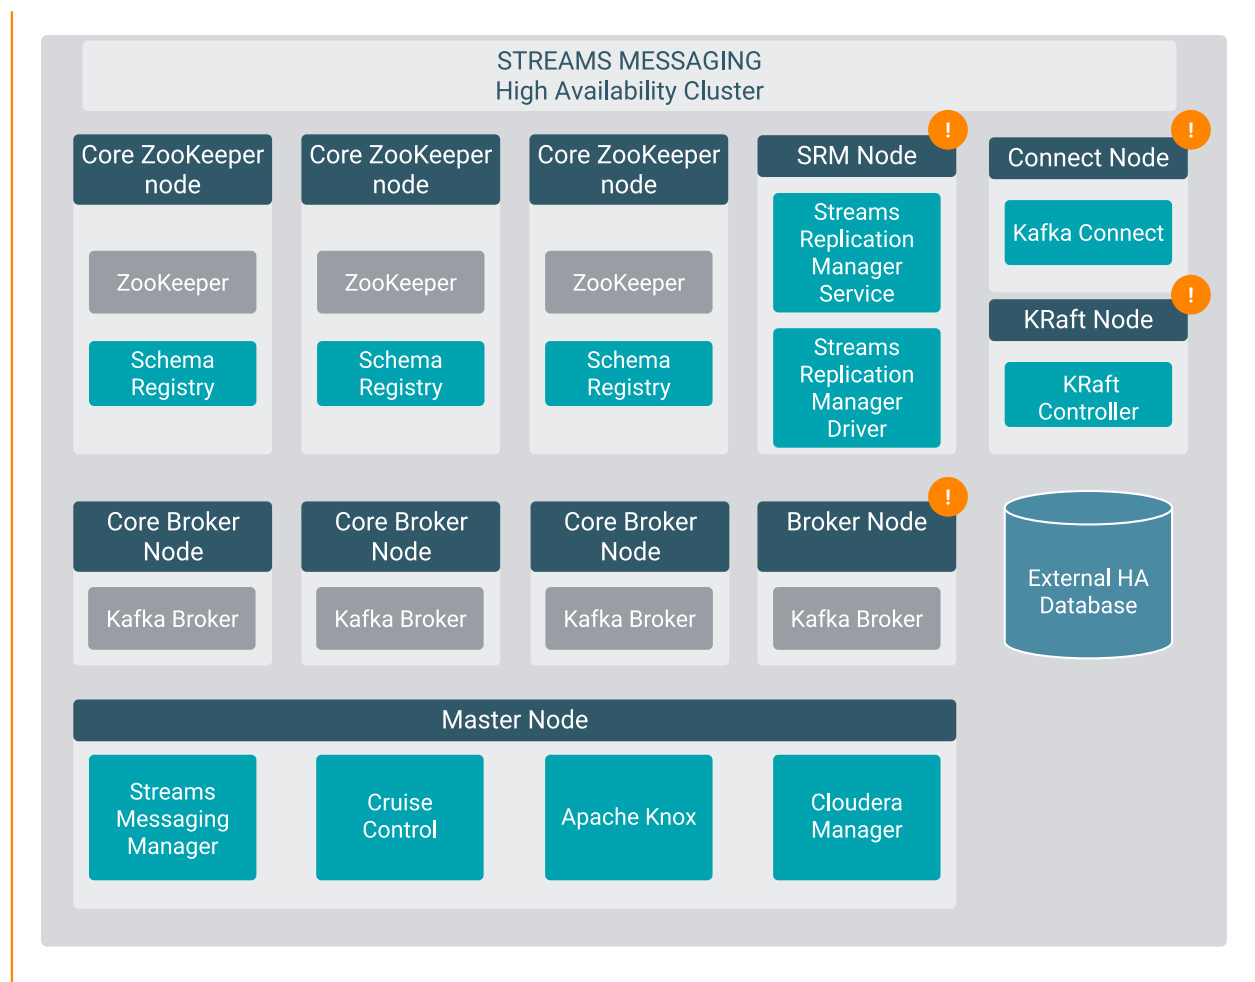
Important: The following is the currently available layout for High Availability clusters. This is the layout of all High Availability clusters provisioned after October 30, 2023. See the **V1** tab to view the layout that was available before October 30, 2023.



For V1



Important: The following is the layout of High Availability clusters provisioned before October 30, 2023. This layout is no longer available for provisioning. See the **V2** tab to view the currently available layout.



Non-default nodes and scaling

SRM, Broker, Connect, and KRaft nodes are not provisioned by default. If you want to have any of these services provisioned, you must manually set the instance count of the appropriate host group to at least one during cluster provisioning. Otherwise, the host group and its nodes are not provisioned. After a cluster is provisioned, you also have the option to scale these nodes. For more information on scaling, see *Scaling Streams Messaging Clusters*.

Use multiple subnets for high availability

When using the Streams Messaging High Availability definition, ensure that you select multiple subnets when provisioning the cluster. Otherwise, your cluster will not be highly available.

KRaft and Zookeeper

By default, Kafka uses Zookeeper as its metadata store. If you provision KRaft nodes in the cluster, Kafka runs in KRaft mode and uses KRaft as its metadata store. In this case, ZooKeeper instances are still provisioned. However, ZooKeeper is not used by Kafka to store and manage metadata. If required, ZooKeeper can be removed from the cluster after the provisioning is finished. ZooKeeper can be removed in Cloudera Manager. For more information, see *Deleting ZooKeeper from Streams Messaging clusters*.

You can only provision an odd number of KRaft nodes. This is required so that KRaft can hold a majority election for leadership. While it is possible to run Kafka in KRaft mode with a single KRaft node, Cloudera recommends that you provision a minimum of three to avoid having a single point of failure. Cluster provisioning fails if an even number of KRaft nodes are provisioned.

Broker and Core Broker volume per instance count

By default, the volume per instance count for Broker and Core Broker nodes is identical. If you customize your cluster during provisioning, Cloudera recommends that Attached Volume per Instances is set to the same value for both node types. Alternatively, if you want to provision a cluster where the number of volumes is not identical, ensure that you complete *Configure data directories for clusters with custom disk configurations* after the cluster is provisioned. Otherwise, Kafka does not utilize all available volumes. Additionally, scaling the cluster might not be possible.

Default instance and storage configuration

The following table collects the default instance type and storage configuration of the various nodes deployed with the Streams Messaging: High Availability cluster. For more information about the cloud provider-specific instance and storage types, see the *Related Information* section.

For Azure

Table 7: Streams Messaging: High Availability default hardware configuration in Azure

Node	Instance type	Storage configuration
Master	Standard_E16s_v3	100 GB Standard Locally-redundant SSD storage
Manager ¹	Standard_D16_v3	100 GB Locally-redundant storage
Core Broker	Standard_D8s_v3	1 TB Premium locally-redundant storage
Broker	Standard_D8s_v3	1 TB Premium locally-redundant storage
Core ZooKeeper	Standard_D8s_v3	100 GB Locally-redundant storage
SRM	Standard_D8_v3	100 GB Locally-redundant storage
Connect	Standard_D8_v3	100 GB Locally-redundant storage
KRaft	Standard_D8_v3	100 GB Locally-redundant storage

For AWS

Table 8: Streams Messaging: High Availability default hardware configuration in AWS

Node	Instance type	Storage configuration
Master	r5.4xlarge	100 GB Magnetic
Manager ¹	m5.4xlarge	100 GB Magnetic
Core Broker	m5.2xlarge	1 TB General Purpose (GP3 SSD)
Broker	m5.2xlarge	1 TB General Purpose (GP3 SSD)
Core ZooKeeper	m5.2xlarge	100 GB Magnetic
SRM	m5.2xlarge	100 GB Magnetic
Connect	m5.2xlarge	100 GB Magnetic
KRaft	m5.2xlarge	100 GB Magnetic

For GCP

¹ The Manager node is only available on clusters provisioned after October 30, 2023.

Table 9: Streams Messaging: High Availability default hardware configuration in GCP

Node	Instance type	Storage configuration
Master	e2-highmem-16	100 GB Standard persistent disks (HDD)
Manager ¹	e2-standard-16	100 GB Standard persistent disks (HDD)
Core Broker	e2-standard-8	1 TB Solid-state persistent disks (SSD)
Broker	e2-standard-8	1 TB Solid-state persistent disks (SSD)
Core ZooKeeper	e2-standard-8	100 GB Standard persistent disks (HDD)
SRM	e2-standard-8	100 GB Standard persistent disks (HDD)
Connect	e2-standard-8	100 GB Standard persistent disks (HDD)
KRaft	e2-standard-8	100 GB Standard persistent disks (HDD)

Related Information[Scaling Streams Messaging Clusters](#)[AWS instance types](#)[Azure instance types](#)[GCP instance types](#)[AWS storage information](#)[Azure storage information](#)[GCP storage information](#)[Configure data directories for clusters with custom disk configurations](#)[Deleting ZooKeeper from Streams Messaging clusters](#)