

Private Cluster Support (Preview)

Date published: 2022-01-06

Date modified: 2022-03-11

Legal Notice

© Cloudera Inc. 2022. 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.

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.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Contents

Private Cluster Support	4
Requirements	4
Enable the Private Cluster	4
Notes	5

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Private Cluster Support

Each type of network architecture supported by CDP has a unique set of tradeoffs among ease of setup, security, workloads (Experiences) supported, and so on.

Private Clusters provide a simple way to create a secure cluster, where the API server and the workloads themselves only rely on private IP addresses that are not accessible from the internet. Connectivity to the cluster from the CDP control plane is provided by the Cluster Connectivity Manager v2 (CCM v2). CCMv2 uses an agent running in the cluster, and an inverting proxy running on CDP, which creates a HTTPS tunnel between the workload and the control plane.

Requirements

Ensure the following entitlements are needed to enable provisioning of private Kubernetes clusters. Customers should file SRE Jira tickets to request these entitlements, if necessary.

- CDP_CCM_V2_JUMPGATE
- CCMV2 JUMPGATE
- ML_ENABLE_PRIVATE_CLUSTER

Enable the Private Cluster

To enable a private cluster, select the option when provisioning the workspace.

1. In **ML Workspaces**, select **Provision Workspace**.
2. Enter a **Workspace Name**, and select **Environment**.
3. Select the **Advanced Options** toggle.
4. In **Network Settings**, select **Enable Fully Private Cluster**.
5. Make any other settings needed, and select **Provision Workspace**.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Network Settings

Subnets ⓘ

Select Subnets

Load Balancer Source Ranges ⓘ

0.0.0.0/0

Enable Fully Private Cluster

Enable Public IP Address for Load Balancer

The workspace is provisioned using a fully private cluster.

Notes

- In a private cluster configuration, the Kubernetes API server has a private IP address, and it is not routable from the internet, thus reducing the attack surface for the cluster. If a customer wants to have internet access, they need to peer their on-premises network to the CDP environment through VPN peering, for example. The API server is not restricted to only having a private IP address. There are some cases such as with Azure, where for example the load balancer has a public IP address provisioned by Microsoft.
- An alternative configuration that is supported is to provide a non-transparent proxy server for the private cluster (for AWS only). For more information, see <https://docs.cloudera.com/machine-learning/cloud/requirements-aws/topics/ml-non-transparent-proxy-aws.html>

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.