

Azure VM Encryption at Host (Preview)

Date published: 2022-06-06

Date modified: 2022-06-06

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.

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.

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

Introduction	4
Limitations	4
Prerequisites	4
Enable encryption at host for an environment	5
Enable encryption at host for a Data Hub	8

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.

Introduction

You can optionally enable encryption at host for Data Lake, FreeIPA, and Data Hubs. Currently, you need to enable it individually for each Virtual Machine (VM) on Azure Portal.

As described in [Azure documentation](#), when you enable encryption at host, the encryption starts on the VM host, where the data for your temporary disk, and OS and data disk caches are stored. After enabling encryption at host, all this data is encrypted at rest and flows encrypted to the Storage service, where it is persisted. Thus, encryption at host essentially encrypts your data from end to end.

Temporary disks and ephemeral OS disks are encrypted at rest with platform-managed keys when you enable end-to-end encryption. The OS and data disk caches are encrypted at rest with either customer-managed keys (CMK) or platform-managed keys, depending on the selected disk encryption type. For example, if a disk is encrypted with customer-managed keys, then the cache for the disk is encrypted with customer-managed keys, and if a disk is encrypted with platform-managed keys then the cache for the disk is encrypted with platform-managed keys.

Encryption at host does not use your VM's CPU and doesn't impact your VM's performance.

Related links

[Encryption at host - End-to-end encryption for your VM data](#)

Limitations

When using Azure VM encryption at host with CDP, the following limitations apply:

- Even if you wish to use the feature for a single subscription only, you need to enable encryption at host for all subscriptions within your Azure tenant.
- This feature can currently be configured individually for each Data Lake and FreeIPA node. Additionally, encryption at host for Data Hub nodes can be configured individually for each Data Hub node. The configuration must be performed on the Azure Portal.

Prerequisites

Prior to enabling encryption at host in CDP, meet the following Azure prerequisites:

1. Enable encryption at host, as described in [Use the Azure CLI to enable end-to-end encryption using encryption at host: Prerequisites](#) in Azure documentation.

Note: You need to enable this for each subscription within your Azure tenant.

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.

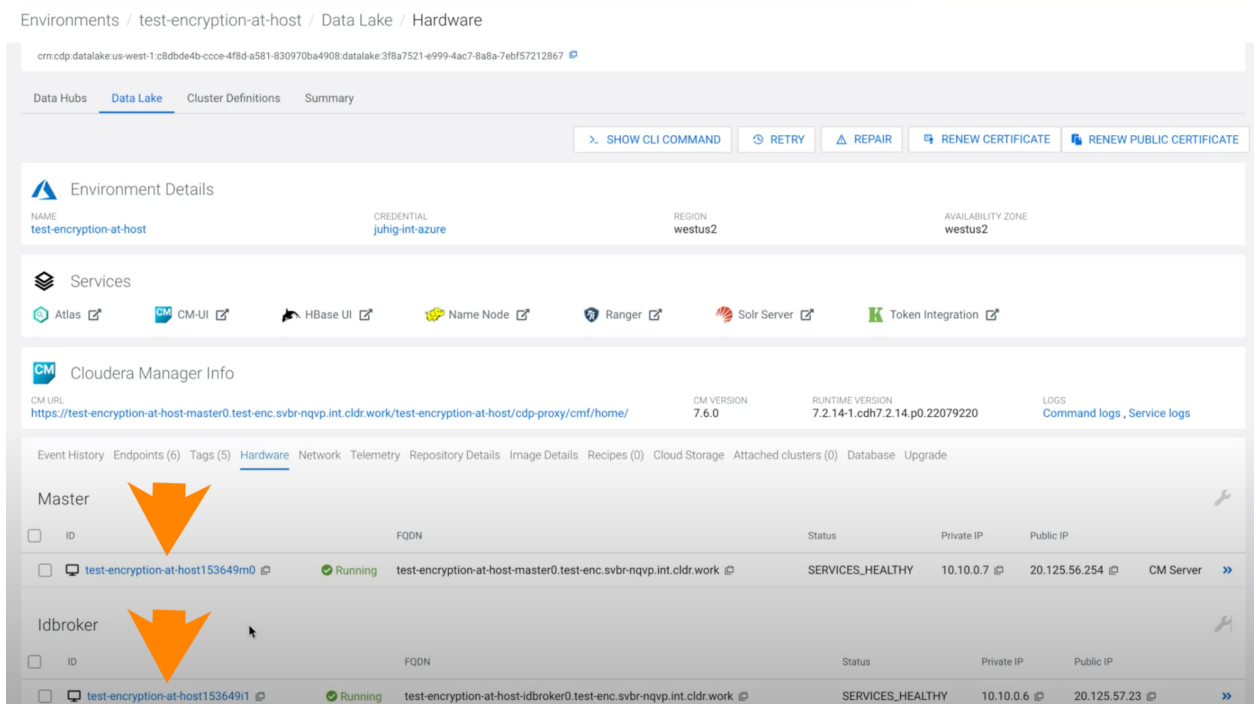
2. If you would like to use Azure disk encryption with a customer-managed key (CMK) along with encryption at host, meet the prerequisites mentioned in [Customer managed encryption keys](#).

Enable encryption at host for an environment

Use these steps to enable encryption at host for an existing CDP environment running on Azure. The steps involve manually enabling encryption at host for each Data Lake and FreeIPA VM via the Azure Portal.

Steps

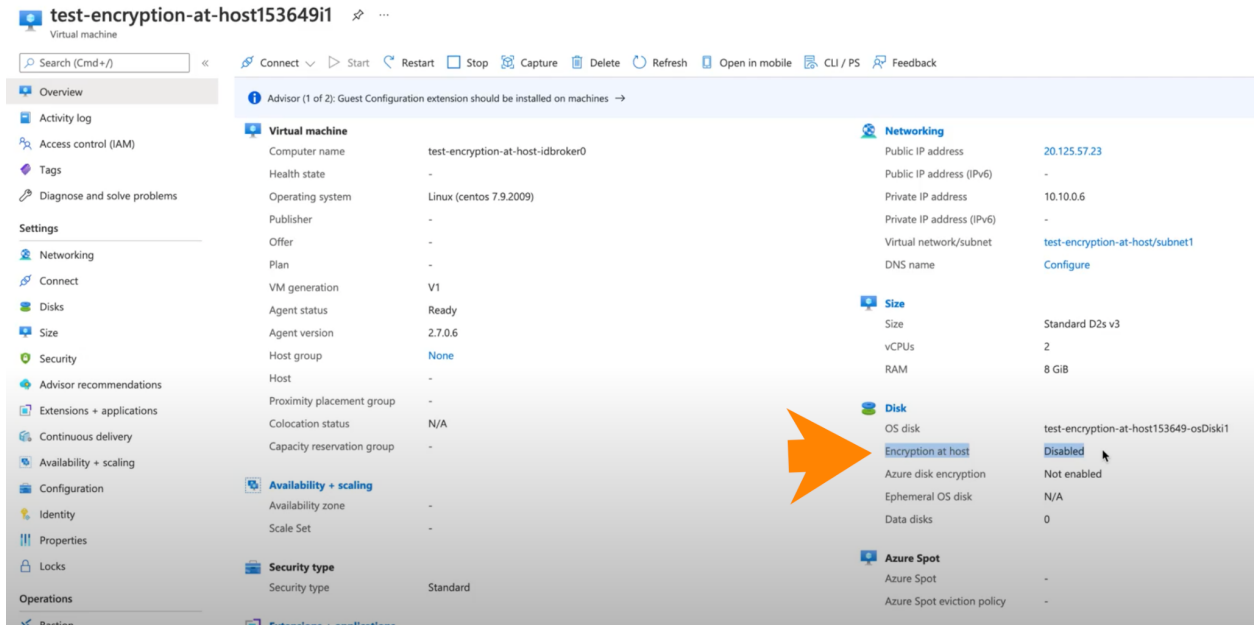
1. In the CDP Management Console, select **Environments** and then click on the specific environment.
2. Make sure that your CDP environment is running.
3. Click on the **Data Lake** tab and then navigate to the **Hardware** tab. Here you can find the list of all Data Lake VMs organized into host groups:



4. Click on each of the VM links and a new browser tab will open for each, redirecting you to the Azure Portal. You need to do this individually for each VM.
5. For each of the VMs, navigate to the **Disks** section in Azure Portal. It will show the “Encryption at Host” as “disabled”:

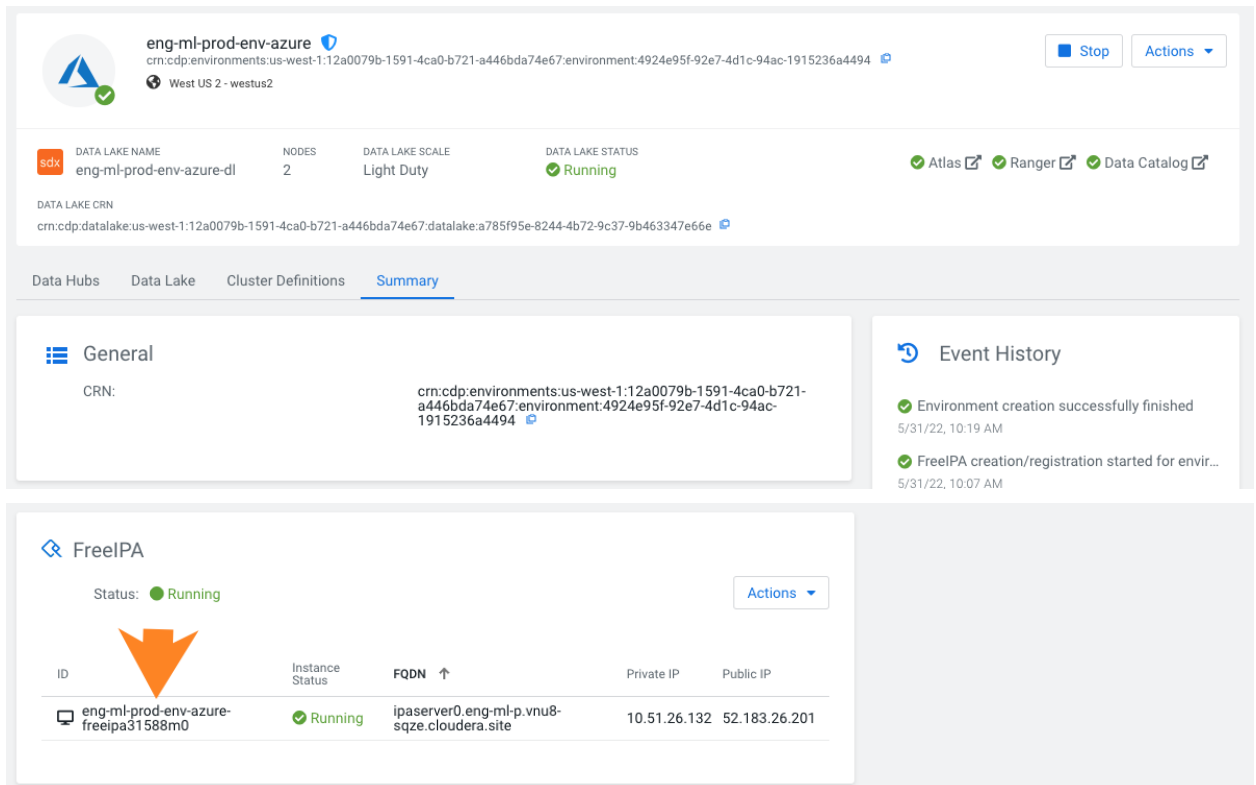
This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes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

CLUDERA TECHNICAL PREVIEW DOCUMENTATION



Leave the Azure Portal browser tabs open. You will need to get back to them shortly.

6. Navigate back to the CDP Management Console to repeat the same steps for FreeIPA VMs. To access FreeIPA VMs, navigate to environment details and click on the **Summary** tab. Next, scroll down to the **FreeIPA** tile. Here you can find the list of all FreeIPA VMs organized into host groups:

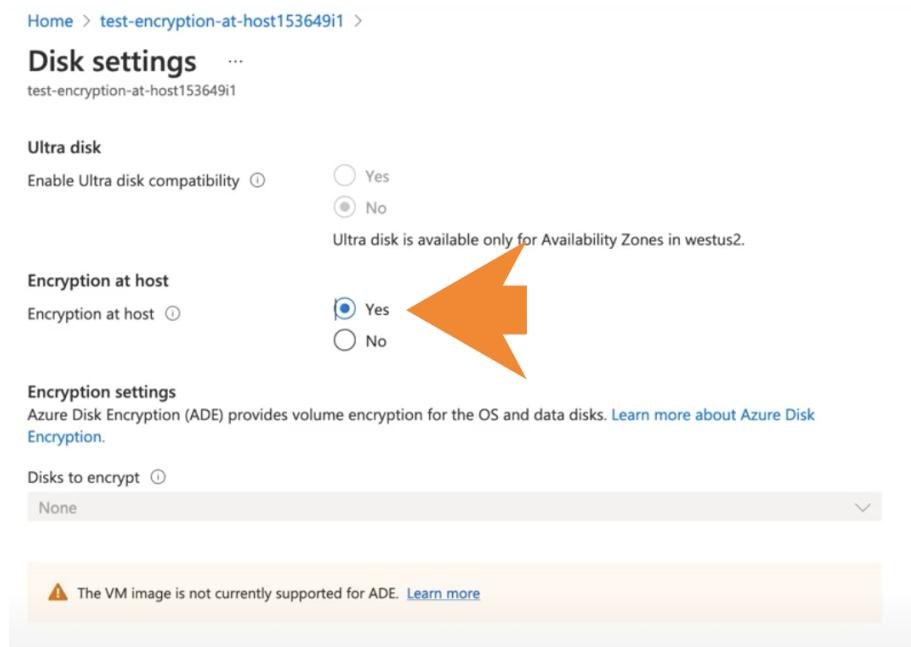


There is no link to Azure Portal, but you can copy the IDs of the VMs and search for them on Azure Portal. Just as you did for each Data Lake node, for each FreeIPA node,

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.

navigate to the **Disks** section in Azure Portal. It will show the “Encryption at Host” as “disabled”. Again, leave the Azure Portal browser tabs open. You will need to get back to them shortly.

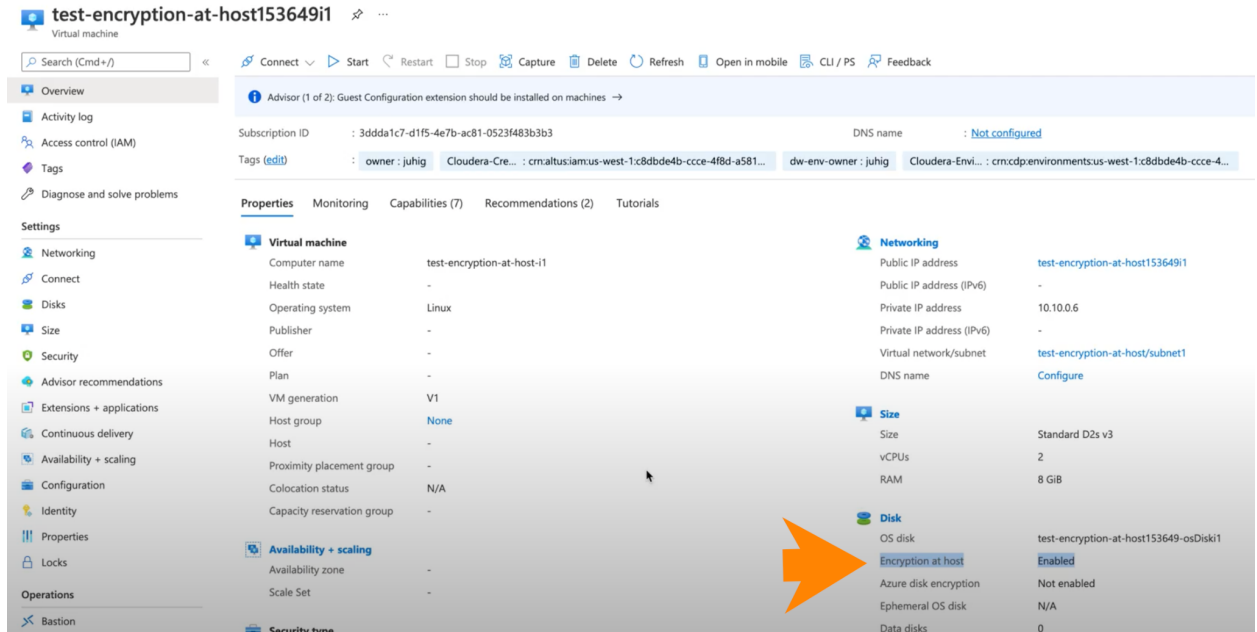
7. Navigate back to the CDP Management Console and stop the environment by clicking the **Stop** button in the top right corner in environment details. If you need detailed instructions, see [Stop and restart an environment](#).
8. Once the environment has been successfully stopped, navigate back to the Azure Portal browser tabs opened earlier.
9. In Azure Portal, perform this for each VM of the Data Lake and for each VM of FreeIPA:
 - a. Navigate to the **Disks** section.
 - b. Within the Disks tab, navigate to the **Additional settings** section.
 - c. Select “Yes” for the “**Encryption at Host**” setting:



- d. Click on **Save**.
 - e. Once the update is complete, you will see a message stating “Updated virtual machine”.
10. Before proceeding, ensure that you have performed the above steps for all Data Lake VMs and for all FreeIPA VMs.
11. Navigate back to the browser tab with the CDP Management Console and restart the environment by clicking the **Start** button in the top right corner in environment details. If you need detailed instructions, see [Stop and restart an environment](#).
12. Once the environment has been successfully restarted, find the **Hardware** section in the **Data Lake** tab, just like you did earlier, and click on each of the Data Lake VM links. A new browser tab will open for each, redirecting you to the Azure Portal. For each of these VMs, navigate to the **Disks** section in Azure Portal. It will show the “Encryption at Host” as “enabled”:

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.

CLUDERA TECHNICAL PREVIEW DOCUMENTATION



Next, confirm the same for all FreeIPA VMs in the **Summary** tab > **FreeIPA** tile.

Enable encryption at host for a Data Hub

Use these steps to enable encryption at host for an existing Data Hub running on Azure. The steps involve manually enabling encryption at host for each Data Hub VM via the Azure Portal.

Before you begin

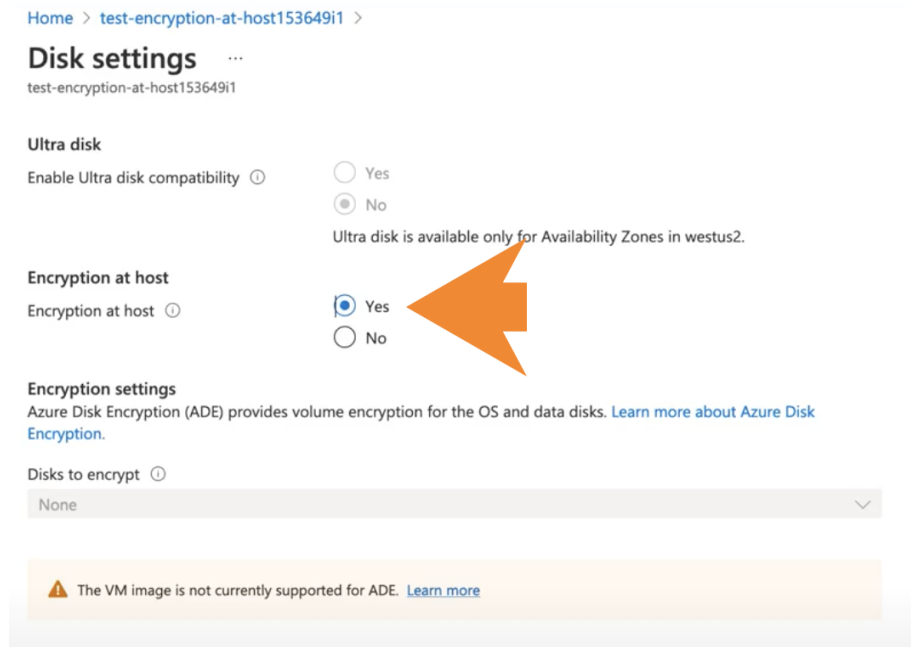
Not all Azure VM types support encryption at host. In order to use encryption at host, when creating your Data Hub, select VM types that support encryption at host for each Data Hub host group. VM types can be selected per host group during Data Hub creation in **Advanced Options > Hardware and Storage**. To find which VM types support encryption at host, follow the steps in [Finding supported VM sizes](#).

Steps

1. In the CDP Management Console, select **Data Hubs**, and click on the specific Data Hub.
2. Make sure that your Data Hub is running.
3. In Data Hub details, navigate to the **Hardware** tab. Here you can find the list of all Data Hub VMs organized into host groups:

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.

8. Once the Data Hub has been successfully stopped, navigate back to the Azure Portal browser tabs opened earlier.
9. In Azure Portal, perform the following for each Data Hub VM:
 - a. Navigate to the **Disks** section.
 - b. Within the Disks tab, navigate to the **Additional settings** section.
 - c. Select “Yes” for the “**Encryption at Host**” setting:



- d. Click on **Save**.
 - e. Once the update is complete, you will see a message stating “Updated virtual machine”.
10. Before proceeding, ensure that you have performed the above steps for all Data Hub VMs.
11. Navigate back to the browser tab with the CDP Management Console and restart the Data Hub by clicking the **Start** button in the top right corner. If you need detailed instructions, see [Restart a cluster](#).
12. Once the environment has been successfully restarted, find the **Hardware** section in the Data Hub details, just like you did earlier, and click on each of the Data Hub VM links.. A new browser tab will open for each, redirecting you to the Azure Portal. For each of these VMs, navigate to the **Disks** section in Azure Portal. It will show the “Encryption at Host” as “enabled”:

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.

CLOUDERA TECHNICAL PREVIEW DOCUMENTATION

The screenshot displays the Azure portal interface for a virtual machine. The left sidebar contains navigation options such as Overview, Activity log, Access control (IAM), Tags, Diagnose and solve problems, Settings, Networking, Connect, Disks, Size, Security, Advisor recommendations, Extensions + applications, Continuous delivery, Availability + scaling, Configuration, Identity, Properties, Locks, Operations, and Bastion. The main content area shows the VM's properties, including Subscription ID, DNS name, and Tags. The 'Properties' tab is active, displaying details for the 'Virtual machine' and 'Availability + scaling' sections. The 'Disk' section is highlighted with an orange arrow, showing the following configuration:

Property	Value
OS disk	test-encryption-at-host153649-osDisk1
Encryption at host	Enabled
Azure disk encryption	Not enabled
Ephemeral OS disk	N/A
Data disks	0

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.