

Hitachi Content Intelligence

Getting Started Guide

Legal notice

© 2016, 2022 Hitachi Vantara Corporation. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara Corporation (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara Corporation at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara Corporation.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, IBM, Lotus, MVS, OS/390, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found at <https://www.hitachivantara.com/en-us/company/legal.html>.

Preface

This book is the installation guide for HCI. It describes how to install both multi and single-instance HCI systems that power either the HCM or Hitachi Content Search use cases.

Please read this document carefully to understand how to install these products, and maintain a copy for your reference.

About this document

This Getting Started Guide is intended to provide guidance for installing a single-instance HCM deployment with a basic configuration. For more complex deployments and additional information, see *Installing HCI* bundled with your product.

Intended audience

This document contains information on installing an HCI system, either on physical servers or virtual machines that you provide.

Product version

This document applies to Hitachi Content Intelligence 1.6 or later.

Release notes


Read the release notes before installing and using this product. They may contain requirements or restrictions that are not fully described in this document or updates or corrections to this document. Release notes are available on the Hitachi Vantara Support Website: <https://knowledge.hitachivantara.com/Documents>.






Document conventions

This document uses the following typographic conventions:

Convention	Description
Bold	<ul style="list-style-type: none"> Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click OK. Indicates emphasized words in list items.
<i>Italic</i>	<ul style="list-style-type: none"> Indicates a document title or emphasized words in text. Indicates a variable, which is a placeholder for actual text provided by the user or for output by the system. Example: <pre>pairdisplay -g group</pre> <p>(For exceptions to this convention for variables, see the entry for angle brackets.)</p>
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairdisplay -g oradb</code>
< > angle brackets	Indicates variables in the following scenarios: <ul style="list-style-type: none"> Variables are not clearly separated from the surrounding text or from other variables. Example: <pre>Status-<report-name><file-version>.csv</pre> Variables in headings.
[] square brackets	Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a b } indicates that you must choose either a or b.
vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b.

This document uses the following icons to draw attention to information:

Icon	Label	Description
	Note	Calls attention to additional information.

Icon	Label	Description
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Important	Highlights information that is essential to the completion of a task.
	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).
	CAUTION	Warns the user of a hazardous situation that, if not avoided, could result in major or minor injury.
	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

Conventions for storage capacity values

Physical storage capacity values (for example, disk drive capacity) are calculated based on the following values:

Physical capacity unit	Value
1 kilobyte (KB)	1,000 (10^3) bytes
1 megabyte (MB)	1,000 KB or $1,000^2$ bytes
1 gigabyte (GB)	1,000 MB or $1,000^3$ bytes
1 terabyte (TB)	1,000 GB or $1,000^4$ bytes
1 petabyte (PB)	1,000 TB or $1,000^5$ bytes
1 exabyte (EB)	1,000 PB or $1,000^6$ bytes

Logical capacity values (for example, logical device capacity, cache memory capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 cylinder	Mainframe: 870 KB

Logical capacity unit	Value
	Open-systems: <ul style="list-style-type: none"> ▪ OPEN-V: 960 KB ▪ Others: 720 KB
1 KB	1,024 (2 ¹⁰) bytes
1 MB	1,024 KB or 1,024 ² bytes
1 GB	1,024 MB or 1,024 ³ bytes
1 TB	1,024 GB or 1,024 ⁴ bytes
1 PB	1,024 TB or 1,024 ⁵ bytes
1 EB	1,024 PB or 1,024 ⁶ bytes

Accessing product documentation

Product user documentation is available on Hitachi Vantara Support Connect: <https://knowledge.hitachivantara.com/Documents>. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

The [Hitachi Vantara Support Website](https://support.hitachivantara.com/en_us/contact-us.html) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to the Hitachi Vantara Support Website for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

[Hitachi Vantara Community](https://community.hitachivantara.com) is a global online community for Hitachi Vantara customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hitachivantara.com, register, and complete your profile.

Comments

Please send comments to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara Corporation.

Thank you!

Contents

Preface.....	3
Intended audience.....	3
Product version.....	3
Release notes.....	3
Document conventions.....	3
Conventions for storage capacity values.....	5
Accessing product documentation.....	6
Getting help.....	6
Comments.....	6
Chapter 1: Planning and preparation.....	8
Items you need.....	8
Chapter 2: Environment setup.....	9
Install Docker on the server or virtual machine	9
Configure Docker on the server or virtual machine.....	9
Configure the mmap count.....	10
Configure the firewall rules on the server or virtual machine.....	10
Run Docker on the server or virtual machine.....	10
Chapter 3: Software installation.....	11
Unpack the installation package.....	11
Start the application on your server or virtual machine.....	12
Chapter 4: HCI deployment.....	13
Access deployment wizard.....	13
Deploy the system.....	14

Chapter 1: Planning and preparation

This section of the HCI Getting Started guide helps you prepare for installation of HCI.

Items you need

To install a system, you need the `HCI-<version-number>.tgz` file.

This archive file includes the software installation files needed to install your HCI instance.

Chapter 2: Environment setup

This section of the HCI Getting Started guide helps you prepare your environment for the HCI product.

Install Docker on the server or virtual machine

Procedure

1. In a terminal window, check whether Docker 1.13.1 or later is installed:
`docker -version`
2. If Docker is not installed or if you have a version prior to 1.13.1, install the current Docker version suggested by your operating system.

The installation method you use depends on your operating system. See the [Docker website](#) for instructions.

Configure Docker on the server or virtual machine

Before installing HCI, configure Docker with settings suitable for your environment. For guidance on configuring and running Docker, see the applicable Docker documentation.

Procedure

1. Ensure that the Docker installation directory on the instance has at least 20 GB available for storing the HCI Docker images.
2. Ensure that the Docker storage driver is configured correctly on the instance.
After installation, changing the Docker storage driver requires reinstallation of HCI.
To view the current Docker storage driver on an instance, run: `docker info`.
3. If you want to enable SELinux on the system instances, use a Docker storage driver that supports it.
The storage drivers that SELinux supports differ depending on the Linux distribution you're using. For more information, see the Docker documentation.
4. If you are using the Docker `devicemapper` storage driver, ensure that there's at least 40 GB of Docker metadata storage space available on the instance. HCI requires 20 GB to install successfully and an additional 20 GB to successfully update to a later version.
To view Docker metadata storage usage on an instance, run: `docker info`



Note: On a production system, do not run `devicemapper` in `loop-lvm` mode. This can cause slow performance or, on certain Linux distributions, HCI may not have enough space to run.

Configure the mmap count

To configure the mmap count, add the following line to the `/etc/sysctl.conf` file:

```
vm.max_map_count = 262144
```

If the line already exists, ensure that the value is greater than or equal to 262144.

Configure the firewall rules on the server or virtual machine

Procedure

1. Edit the firewall rules to allow communication over all network ports that you want your system to use. You do this using a firewall management tool such as `firewalld`.
2. Restart the server or virtual machine.

Run Docker on the server or virtual machine

To unpack your installation package, Docker needs to be both started and kept running. You can use whatever tools you typically use for keeping services running in your environment.

For example, to run Docker using `systemd`:

Procedure

1. Verify that Docker is running:

```
systemctl status docker
```
2. If Docker is not running, start the `docker` service:

```
sudo systemctl start docker
```
3. (Optional) Configure the Docker service to start automatically when you restart the server or virtual machine:

```
sudo systemctl enable docker
```

Chapter 3: Software installation

This section of the HCI Getting Started guide helps you with the software installation of the HCI product.

Unpack the installation package

1. Download the installation package `HCI-<version_number>.tgz` and store it in a directory on the server or virtual machine.
2. In the largest disk partition on the server or virtual machine, create a directory named `/hci`:

```
mkdir /<path>/hci
```

3. Move the installation package to the directory `/hci`:

```
mv /<path>/HCI-<version_number>.tgz /<path>/hci
```

4. Navigate to the installation directory:

```
cd /<path>/hci
```

5. Unpack the installation package:

```
tar -xzf HCI-<version_number>.tgz
```

This creates a number of directories and folders within the installation directory.

6. Run the installation script `install`, located within the directory matching the version number of the installation package:

```
sudo ./cluster/<version_number>/bin/install
```



Note:

- Don't change directories after running the installation script. The following steps are performed in your current directory, not under the version-specific directory.
- The installation script can be run only once on an instance. You cannot rerun this script to try to repair or upgrade a system instance.

Start the application on your server or virtual machine

Procedure

1. Start the application script `run` using whatever methods you usually use to run scripts.



Important: Ensure that the method you use can keep the `run` script running and can automatically restart it in case of a server reboot or other availability event.

Here are some examples of how you can start the script:

- You can run the script in the foreground:

```
sudo /<path>/hci/bin/run
```

When you run the `run` script this way, the script does not automatically complete, but instead remains running in the foreground.

- You can run the script as a service using `systemd`:

- a. Copy the `hci.service` file to the appropriate location for your OS. For example:

```
cp /<path>/hci/bin/hci.service /etc/systemd/system
```

- b. Enable and start the `hci.service` service:

```
sudo systemctl enable hci.service
sudo systemctl start hci.service
```



Note: When you enable the search service, `systemctl` may display this message:

The unit files have no [Install] section. They are not meant to be enabled using `systemctl`. Possible reasons for having this kind of units are:

- 1) A unit may be statically enabled by being symlinked from another unit's `.wants/` or `.requires/` directory.
- 2) A unit's purpose may be to act as a helper for some other unit which has a requirement dependency on it.
- 3) A unit may be started when needed via activation (socket, path, timer, D-Bus, udev, scripted `systemctl` call, ...).

Depending on your OS, the search service may or may not have successfully been enabled.

To avoid this, make sure that you moved the `HCI.service` to the appropriate location, typically `/etc/systemd/system`.

Chapter 4: HCI deployment

This section of the HCI Getting Started guide helps you with the service deployment wizard of the HCI product.

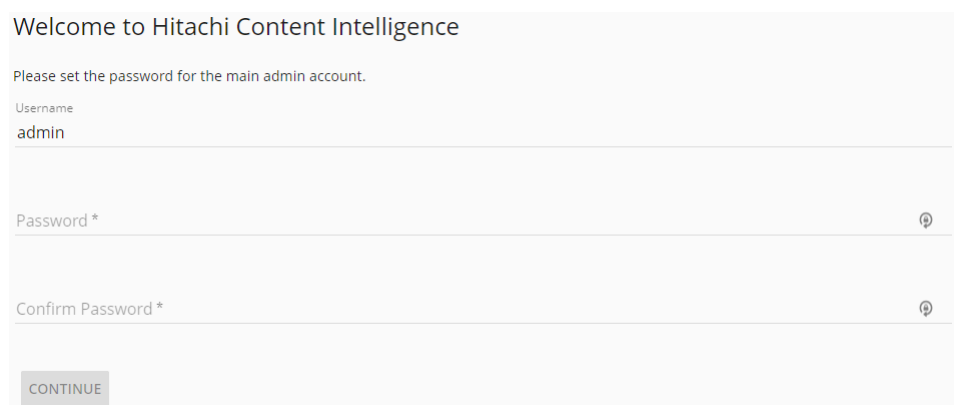
Access deployment wizard

After creating your instance, you need to go to the service deployment wizard in the Admin App.

To access the service deployment wizard:

Procedure

1. Open a web browser and go to: `https://instance_ip_address:8000`
2. On the **Welcome** page, set a password for the admin user account. Then click **CONTINUE**.



Welcome to Hitachi Content Intelligence

Please set the password for the main admin account.

Username
admin

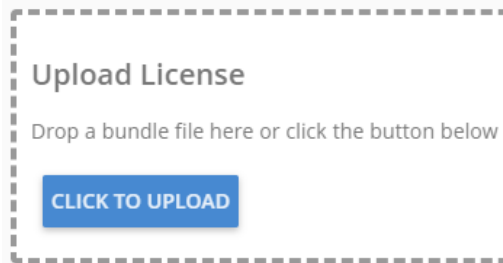
Password *

Confirm Password *

CONTINUE

3. On the **Licensing** page:

- If you have your purchased license file, drag and drop it into the **Upload License** section. Alternatively, you can use the **CLICK TO UPLOAD** button to browse for your license file.



- If you've purchased a license but have not yet received it, make note of the value in the System ID section on the **Licensing** page and contact your sales representative.
- To use the system for a limited amount of time with the pre-installed trial license, click **Next**.

If for some reason the trial license failed to install, there is a copy included in the HCI-
<version-number>.tgz installation package that you can upload to the **Licensing** page.
The trial license is located in the installation package at:

/product/<version>/trial-<version>.plk

4. On the **Set Cluster Hostname/IP** page, specify the hostname for your system. Omitting this can cause links in the Admin App to function incorrectly. Then click **CONTINUE**.

Set the hostname or IP Address

The address entered below will be used to access Hitachi Content Intelligence and all associated apps.

Cluster Hostname/IP Address

cluster119bi.lab.archivas.com

Cluster Hostname or Primary IP Address used to access the apps in this system.

CONTINUE

BACK

5. On the **Choose Deployment** page, select the HCI deployment type that you purchased, either Hitachi Content Search or Hitachi Content Monitor. Then click **CONTINUE**. The **Confirm Cluster Topology** page displays and shows the detected instance.



Note: If the instance is not displayed, click the **Click here** link in the **Instance Discovery** panel until they appear.

Deploy the system

When all your instances have been discovered and configured, click **Deploy Single Instance** or **Deploy Cluster**. The system deployment starts. You can click the **View Deployment Details** link to view the process of the deployment.

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
HitachiVantara.com | community.HitachiVantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact