

Hitachi Content Intelligence

v2.1

Release Notes

This release of Hitachi Content Intelligence (HCI) includes new HCP for Cloud Scale S3 functionality, an upgrade path for previous versions of the product to Solr 8, and multiple bug and security fixes.

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 LLC (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 LLC 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 LLC.

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, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, 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>.

Contents

Hitachi Content Intelligence v2.1 Release Notes.....	4
About this document.....	4
Intended audience.....	4
HCI Community.....	4
About this release.....	4
Major features.....	5
General updates.....	5
Resolved issues.....	6
Known issues.....	7
System requirements.....	10
Hardware requirements.....	10
Software requirements.....	10
Operating system and Docker minimum requirements.....	11
Operating system and Docker qualified versions.....	11
Docker considerations.....	12
SELinux considerations.....	12
Time source requirements.....	12
Supported browsers.....	12
Documentation set.....	13

Hitachi Content Intelligence v2.1 Release Notes

About this document

This document provides information about the Hitachi Content Intelligence (HCI) software, including new features, important enhancements, as well as fixed and known issues.

Intended audience

This document is intended for system administrators, Hitachi Vantara representatives, and authorized service providers who configure and operate the HCI software.

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.

HCI Community

For HCI-specific product support, discussions, announcements, and FAQs, visit us at the HCI Community portal: <https://community.hitachivantara.com/s/hitachi-content-intelligence>

About this release

This release of Hitachi Content Intelligence (HCI) includes new HCP for Cloud Scale S3 functionality, an upgrade path for previous versions of the product to Solr 8, and multiple bug and security fixes.

Major features

General updates

Updates and upgrading to Solr 8

HCI v2.0 upgraded Solr from v6 to v8. Because of this, it was a fresh-install only release and did not support updates. HCI v2.1 now features update support for previous versions to gain the new Solr 8 functionality

On an update from HCI v1.x to v2.1, you will now have two Solr services available:

- The existing Index service for Solr 6.
- The new Solr 8 Index service.

This new version of Solr has improved indexing capabilities and will allow you to perform incremental back up and restore of your data. It will remain undeployed and needs to be added by clicking **ADD SERVICE** under the **Services** tab in the Admin App. Once deployed, both Solr 6 *and* Solr 8 services will be running simultaneously on your system.

It is up to the user to reindex all of their data to Solr 8 indexes. Once that is done, the Solr 6 indexes can be deleted from the system and you can eventually scale off the Solr 6 Index service. Alternatively, a fresh install of HCI v2.1 will include *only* the Sol 8 Index service scaled up and fully functional.



Important: As Solr 6 indexes will not be able to take advantage of future HCI security upgrades, their use in HCI will be deprecated. You can still use your existing Solr 6 indexes but it is strongly recommended to use Solr 8 indexes to remain fully protected.

For more information and best practices around reindexing, see the [Apache Solr Reference Guide](#).

New HCP for Cloud Scale functionality

The S3 compatible connector in Workflow Designer now supports object lock through the **Include S3 Object lock metadata** field, which reads both the retention period and legal hold metadata of S3 objects.

The connector now also reads object lock metadata through the **Set Legal Hold** and **Set Retention** actions, as well as object tagging metadata through the **Delete Tags** and **Set Tags** actions.

In addition, the **Add Metadata** and **Attach Stream** stages now allow you to retrieve a specific version of an object, and the **Execute Action** stage allows you to delete a specific version, all based on an object's S3 version ID.

New HCP for Cloud Scale bulk actions

In tandem with the S3 enhancements in Workflow Designer, bulk actions in the Search App have been updated to support the following new S3 actions:

- **S3 Compatible Apply Hold**
- **S3 Compatible Clear Hold**
- **S3 Compatible Delete Object**
- **S3 Compatible Set Retention**
- **S3 Compatible Set Tags**

New package export encryption

As an added layer of security, when exporting packages in the Admin App, users can now include an encryption passphrase to help better protect them. This passphrase is entirely optional, and packages with no passphrase provided will not require one upon import.



Note: Packages exported prior to HCI v2.1 will be able to be imported without any password requirements.

Resolved issues

Issue	Area affected	Description	Outcome
HCI-6104	Monitor App	Several config fields for Anomaly Detection and Forecasting do not have validation and will accept any erroneous values entered.	An error message now appears when users enter a negative number, a non-number, or leave a blank value in the Container Memory and CPU config fields.
HCI-6303	Workflow Designer App	If a workflow test is run while that workflow test is already currently running, a 409 error should be returned, but a 400 error is displayed instead.	The correct 409 error message is now displayed.
HCI-6971	Workflow Designer App	HCI currently checks for and rejects the use of the backslash ("\") character in businessObject names, but there is no check created for the forwardslash ("/"). Using a forwardslash is equally as detrimental and can prevent the importing and execution of workflows.	A check has now been added for the forwardslash ("/") character in businessObject names.

Issue	Area affected	Description	Outcome
HCI-7342	Solr	Incremental backup attempts fail after a shard split operation has completed. This is an internal HCI ticket to track the progress of SOLR-15673, which was filed directly with Apache.	This behavior has been fixed by Apache and the current version of Solr has been implemented in the HCI v2.1 release.
HCI-7363	Solr	After an incremental backup, the <code>purgeUnused</code> operation incorrectly deletes all backed up data. This is an internal HCI ticket to track the progress of SOLR-15706, which was filed directly with Apache.	This behavior has been fixed by Apache and the current version of Solr has been implemented in the HCI v2.1 release.
HCI-7400	Documentation	The "Index restore API" topic in the Workflow Designer Help contains several inaccuracies and needs to be corrected.	The Workflow Designer Help now includes the correct documentation for the "Index restore API" topic.

Known issues

Issue	Area affected	Description	Outcome/Workaround
HCI-286	Workflow Designer App	Using Retry Failed Documents manually in a workflow that has failures doesn't add to the output metrics after it has successfully indexed those files.	The results of the Retry Failed Documents setting do not affect output metrics and the files are still indexed successfully.
HCI-353	Job Driver	When attempting to migrate data to a namespace that has already hit its hard quota, 413 errors are received without any failure notifications in the workflow. The failures are instead reported in the advanced historic metrics.	N/A
HCI-1047	Content Monitor App	Replication links with forward slashes (/) do not appear in the Replication metrics.	Replace the forward slashes in the Replication link name with any other ASCII character (space, %, etc.).
HCI-1737	Workflow Designer	Enabling Process all documents ignores the HCP and Sharepoint connector's directories when a workflow is resumed.	N/A

Issue	Area affected	Description	Outcome/Workaround
HCI-1918	Metrics	Metrics service cannot be scaled to twice the current number of instances.	<ul style="list-style-type: none"> ▪ Scale the instances back down to the original instances. ▪ Wait for the scale action to complete. ▪ Scale the instances back up to the desired number of instances.
HCI-6103	Workflow Designer App	On a multi-node cluster, if one of the nodes goes down, the workflow halts with a task error. The task error states that the driver heap limit is too low for the workflow, but raising this value in the workflow settings does not restart the workflow.	The workflow-agent job type is configured to run on all nodes in a cluster by default. Scaling the workflow-agent off of the bad node will allow workflows to resume.
HCI-6109	Admin App	HCI enters a bad state after rebooting 2 of the 3 master nodes available on a 4-node cluster.	Reboot all HCI master nodes at the same time and the clusters will return normally.
HCI-6128	Metrics	Metrics from historical logs are ignored after an update and present an OOM message.	If you notice your imported logs missing after update, reimport them and the metrics will display correctly.
HCI-6304	Workflow Designer App	Testing a data connection accessible by proxy shows a blank certificate and cancelling it makes the system seem unresponsive.	<p>If your data connection can only be accessed by proxy and requires a certificate, download the certificate from the other system and manually add it to the Admin App</p> <p>To do this, navigate to Configuration > Certificates > Client and click UPLOAD CLIENT CERTIFICATE.</p>
HCI-6540	Solr	During an update, Solr shards appear in the Gone state if Marathon references the old port number and node name pairing when starting the Solr service.	<p>Delete the Gone shards:</p> <ol style="list-style-type: none"> 1. Navigate to the Solr UI. 2. Click Collections. 3. For each affected index: <ol style="list-style-type: none"> a. Click the index name. b. Click the affected shards. c. Click X.

Issue	Area affected	Description	Outcome/Workaround
			d. Confirm the deletion.
HCI-6857	Search App	Search App exclusive users with bulk action permissions are able to see Workflow Designer as an option in their SSO menu. Clicking it presents a message telling the user the page cannot be displayed.	N/A
HCI-7016	Solr	When changing a field type in the schema, a message prompt appears letting the user know that the change will only affect documents which are indexed moving forward. This is incorrect, as changes to the <code>docValues</code> field result with no additional documents being indexed at all.	N/A
HCI-7020	Admin App	When updating from HCI 1.6.x to 1.10.1, the <code>doc</code> folder is owned by <code>root</code> on several nodes in the cluster, causing an update failure.	Changing the user from <code>root</code> to <code>hci</code> and retrying the failure resolves the issue.
HCI-7149	Content Monitor	HCM generates 404 errors on HCP MAPI, which causes them to incorrectly appear in visualizations.	N/A
HCI-7229	Admin App	If a Solr index is in an unhealthy state, where no replicas of any particular shard are showing as a leader, a change (specifically a reduction) to the IPL could potentially cause loss of data.	N/A
HCI-7341	Solr	When attempting to recreate a delete Solr collection with a different initial schema, the old schema is still present in the UI. This is an internal HCI ticket to track the progress of SOLR-15674, which was filed directly with Apache.	N/A
HCI-7366	Search App	The autocomplete functionality of the search bar in Search App does not work for file names in Chinese after reindexing text field changes.	N/A
HCI-7368	Solr	If an index is created without using a SolrCloud connection URL, attempts to create bulk actions using the index will fail.	N/A

Issue	Area affected	Description	Outcome/Workaround
HCI-7369	Solr	When attempting to remove a copy field from an index, the delete fails with a "Collection not found" error and the index is unable to be written to.	N/A
HCI-7370	Solr	Solr can only use comparatives with small numbers and returns 0 results (in error) when using larger ones.	N/A

System requirements

This section lists the hardware, networking, and operating system requirements for running an HCI system with one or more instances.

Hardware requirements

This table shows the minimum and recommended hardware requirements for each instance in an HCI running Hitachi Content Search.

Resource	Minimum	Recommended
RAM	16 GB	32 GB
CPU	4-core	8-core
Available disk space	50 GB	500 GB

Software requirements

The following table shows the minimum requirements and best-practice software configurations for each instance in an HCI system.

Resource	Minimum	Best
IP addresses	(1) static	(2) static

Resource	Minimum	Best
Firewall Port Access	Port 443 for SSL traffic Port 8000 for System Management App GUI Port 8888 for Content Search App GUI	Same
Network Time	IP address of time service (NTP)	Same

Operating system and Docker minimum requirements

Each server or virtual machine you provide must have the following:

- 64-bit Linux distribution
- Docker version 1.13.1 or later installed
- IP and DNS addresses configured

Additionally, you should install all relevant patches on the operating system and perform appropriate security hardening tasks.



Important: Install the current Docker version suggested by your operating system, unless that version is earlier than 1.13.1. The system cannot run with Docker versions before 1.13.1.

Operating system and Docker qualified versions

This table shows the operating systems, as well as the Docker and SELinux configurations, on which this HCI release was qualified. It acts a point of reference for our customers to better share how we operate within our internal environment and does not represent any requirements that need to be followed within your own.

Operating system	Docker version	Docker storage configuration	SELinux setting
CentOS 7.6	Docker 18.03.1-ce	device-mapper	Enforcing
CentOS 8.1	Docker 19.03.13	overlay2	Enforcing
Red Hat Enterprise Linux 8.1	Docker 20.10.14	overlay2	Enforcing
Ubuntu 18.04.4 LTS	Docker 18.03.1-ce	overlay2	Not Installed

Docker considerations

The Docker installation folder on each instance must have at least 20 GB available for storing the HCI Docker images.

Make sure that the Docker storage driver is configured correctly on each instance before installing HCI. To view the current Docker storage driver on an instance, run `docker info`.



Note: After installing, changing the Docker storage driver requires a reinstallation of HCI.

If you are using the Docker `devicemapper` storage driver:

- Make sure that there's at least 40 GB of Docker metadata storage space available on each instance. HCI needs 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`.
- On a production system, do not run `devicemapper` in `loop-lvm` mode. This can cause slow performance or, on certain Linux distributions, HCI might not have enough space to run.

SELinux considerations

You should decide whether you want to run SELinux on system instances and enable or disable it before installing HCI. To enable or disable SELinux on an instance, you must restart the instance. To view whether SELinux is enabled on an instance, run: `sestatus`

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.

Time source requirements

If you are installing a multi-instance system, each instance should run NTP (network time protocol) and use the same external time source. For information, see support.ntp.org.

Supported browsers

The following browsers are qualified for use with HCI software. Other browsers or versions might also work.

- Google Chrome (latest version as of the date of this publication)
- Microsoft Edge (latest version as of the date of this publication)
- Mozilla Firefox (latest version as of the date of this publication)

Documentation set

Along with your release notes, the following guides and documentation comprise the full set of HCI reference documentation:

- *Hitachi Content Intelligence Installing Hitachi Content Intelligence*
- *Hitachi Content Intelligence Deploying the HCI Example OVF*
- *Hitachi Content Intelligence (HCI) Getting Started Guide*
- *Hitachi Content Intelligence Workflow Designer Help*
- *Hitachi Content Intelligence Search Help*
- *Hitachi Content Intelligence Content Monitor Help*
- *Hitachi Content Intelligence Administrator Help*

To learn more, visit the [HCI Knowledge page](#).

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