

Cloudbreak release notes 2

## Cloudbreak Release Notes

**Date of Publish:** 2019-05-28



<https://docs.hortonworks.com/>

# Contents

<b>About this release.....</b>	<b>3</b>
<b>New features.....</b>	<b>3</b>
<b>Technical preview features.....</b>	<b>3</b>
<b>Behavioral changes.....</b>	<b>3</b>
<b>Fixed issues.....</b>	<b>3</b>
<b>Known issues.....</b>	<b>4</b>
<b>Image catalog updates.....</b>	<b>7</b>
<b>Legal information.....</b>	<b>7</b>

## About this release

Cloudbreak 2.9.1 is a general availability release: It is suitable for production.



**Note:**

Features that are not suitable for production are marked as "Technical Preview" or "TP".

## New features

Cloudbreak 2.9.1 is a maintenance release with no new features.

## Technical preview features

Cloudbreak 2.9.1 includes the following technical preview features. These features should only be used for evaluation and are not suitable for production:

Feature	Documentation link
Data lakes	<a href="#">Working with data lakes (TP)</a>
Azure ADLS Gen2	<a href="#">Configuring access to ADLS Gen2</a>
AWS GovCloud	<a href="#">Deploying on AWS GovCloud</a>

## Behavioral changes

Cloudbreak 2.9.1 introduces the following changes in behavior as compared to previous Cloudbreak versions:

Affected item	Description
LDAP configuration for Cloudbreak	Changes were made to make it easier to configure LDAP server's certificate.

## Fixed issues

The following issues were fixed in Cloudbreak 2.9.1:

Issue	Description	Category
CB-1169	Hive metastore storage location is invalid for ADLS Gen2	Usability
CB-1052	Shut down Cloudbreak if there was a migration error at startup	Usability
CB-990	Document LDAPS configuration for Cloudbreak	Documentation
CB-988	Import mounted certificates to truststore at identity start	Usability
CB-723	Getting non terminated clusters query is slow	Usability

Issue	Description	Category
CB-722	Flowlog query by flowId is slow, because extremely high number of IMAGECOPY entries	Usability
CB-705	Http timeout in cb-cli config	Stability
CB-511	Performance improvements for Cloudbreak	Usability
CB-25	Java 11 update	Stability

## Known issues

Cloudbreak 2.9.1 includes the following known issues:

### Known issues: Cloudbreak

Issue	Description	Workaround
BUG-114632	If you started your cluster on Azure with a Cloudbreak version not greater than 2.8.0, 2.7.2 or 2.4.3, then your instances in any 0 or 1 node-sized host group were neither placed in any availability sets nor have rack-info other than 'Default-rack'.	If rack information regarding the related host group is important to you, you should terminate your affected cluster and relaunch it with Cloudbreak 2.9.1.
BUG-116919	When defining network security group rules on Google Cloud, it is possible to specify an incorrect port range such as "5555-3333", causing the cluster deployment to fail with an error similar to:  <pre>Infrastructure creation failed. Reason: Invalid value for field 'resource.allowed[8].ports[0]': '5555-3333'. Second port cannot be smaller than the first port.</pre>	When defining network security group rules on Google Cloud, make sure to define a valid range.

Issue	Description	Workaround
BUG-117004	<p>When defining network security rules during cluster creation on Azure, when ICMP protocol is used, cluster creation fails with an error similar to:</p> <pre data-bbox="656 323 1040 926">Infrastructure creation failed. Reason: Stack provisioning failed, status code InvalidTemplateDeployment, error message The template deployment is not valid according to the validation procedure. See inner errors for details. Please see https://aka.ms/arm-deploy for usage details, details: Security rule has invalid Protocol. Clues provided: Icmp Allowed values: Tcp,Udp</pre>	<p>When defining network security rules during cluster creation on Azure, do not use the ICMP protocol.</p>
BUG-117005	<p>When defining network security rules during cluster creation on Google Cloud via CLI, when ICMP protocol is specified and a port is specified, cluster creation fails with an error similar to:</p> <pre data-bbox="656 1121 1040 1381">Infrastructure creation failed. Reason: Invalid value for field 'resource.allowed[6].ports[0]': '43543'. Ports may only be specified on rules whose protocol is one of [TCP, UDP, SCTP].</pre> <p>This is because when the ICMP protocol is used, no ports should be specified. The UI already enforces this automatically, but with CLI it is possible to specify a port with the ICMP protocol.</p>	<p>When defining network security rules during cluster creation on Google Cloud via CLI, if you would like to define a rule for using the ICMP protocol, do not specify any ports.</p>
BUG-110998	<p>When creating a cluster, the Cloud Storage page in the create cluster wizard includes an option to provide "Path to Ranger Audit Logs for Hive Property" when "Configure Storage Locations" is enabled. This option should only be available for data lakes and not for workload clusters.</p>	<p>Click on "Do not configure".</p>

Issue	Description	Workaround
BUG-99581	<p>The Event History in the Cloudbreak web UI displays the following message:</p> <p>Manual recovery is needed for the following failed nodes: []</p> <p>This message is displayed when Ambari agent doesn't send the heartbeat and Cloudbreak thinks that the host is unhealthy. However, if all services are green and healthy in Ambari web UI, then it is likely that the status displayed by Cloudbreak is incorrect.</p>	<p>If all services are green and healthy in Ambari web UI, then syncing the cluster should fix the problem.</p>
BUG-110999	<p>The auto-import of HDP/HDF images on OpenStack does not work. This means, that in order to start creating HDP or HDF clusters on OpenStack, your OpenStack admin must import these images manually.</p>	<p>Your OpenStack admin must import these images manually by using the instructions in <a href="#">Import HDP and HDF images to OpenStack</a>.</p>
BUG-112787	<p>When a cluster with the same name as specified in CLI JSON already exists, CLI returns:</p> <pre>ERROR: status code: 403, message: Access is denied.</pre>	<p>To avoid this error, pass the cluster name as a parameter with <code>cb cluster create</code> instead of including cluster name in the CLI JSON definition.</p>

**Known issues: HDP**

The known issues described here were discovered when testing Cloudbreak with HDP versions that are used by default in Cloudbreak. For general HDP known issues, refer to HDP release notes published at <https://docs.hortonworks.com/>.

There are no known issues related to HDP.

**Known issues: HDF**

The known issues described here were discovered when testing Cloudbreak with HDF versions that are used by default in Cloudbreak. For general HDF known issues, refer to HDF release notes published at <https://docs.hortonworks.com/>.

Issue	Description	Workaround
BUG-98865	<p>Blueprint configuration parameters are not applied when scaling an HDF cluster. One example that affects all users is that after HDF cluster upscale/downscale the <code>nifi.web.proxy.host</code> blueprint parameter does not get updated to include the new nodes, and as a result the NiFi UI is not reachable from these nodes.</p>	<p>Configuration parameters set in the blueprint are not applied when scaling an HDF cluster. One example that affects all NiFi users is that after HDF cluster upscale the <code>nifi.web.proxy.host</code> parameter does not get updated to include the new hosts, and as a result the NiFi UI is not reachable from these hosts.</p> <p>HOST1-IP:PORT,HOST2-IP:PORT,HOST3-IP:PORT</p>

**Known issues: Data lake**

Issue	Description	Workaround
BUG-109369	Hive does not start on a HDP 2.6 data lake when Kerberos is enabled.	<ol style="list-style-type: none"> <li>Modify /etc/hadoop/&lt;Ambari-version&gt;/0/core-site.xml and /etc/hadoop/conf.backup/core-site.xml by adding the following: <pre>&lt;configuration&gt; &lt;property&gt;    &lt;name&gt;hadoop.security.authentication.name&lt;/name&gt;   &lt;value&gt;kerberos&lt;/value&gt; &lt;/property&gt; &lt;/configuration&gt;</pre> </li> <li>Restart affected services.</li> </ol>
BUG-116913, BUG-114150	HiveServer2 does not start on an HDP 3.1 cluster attached to a data lake. The following error is printed to Ambari logs: <pre>ERROR client.ServiceClient: Error on destroy 'llap0': not found. Failed: org.apache.hadoop.security.AccessControlException: /user/hive/.yarn/ package/LLAP (is not a directory)</pre>	<ol style="list-style-type: none"> <li>Delete the "/user/hive/.yarn/package/LLAP" file, and then create a new directory in this location with the relevant permissions for the hive user.</li> <li>Start HiveServer2.</li> </ol>

## Image catalog updates

This section lists image catalog updates published for Cloudbreak 2.9.1.

**Default HDP, HDF, and Ambari versions provided with Cloudbreak 2.9.1:**

Platform version	Stack version	Ambari version
HDP 2.6	HDP 2.6.5.0-292	Ambari 2.6.2.2
HDP 3.1	HDP 3.1.0.0-78	Ambari 2.7.3.0
HDF 3.3	HDF 3.3.1.0-10	Ambari 2.7.3.0

## Legal information

Cloudbreak 2.9.1

This software includes copyrightable material owned by Hortonworks, Inc.

Cloudbreak incorporates software from various open source projects released primarily under the Apache Software License 2.0 (“ASLv2”). Other software included may be released under the terms of alternative ASLv2 compatible open source licenses. Please review the license and notice files accompanying the software for additional licensing and copyright information.

For additional information on Hortonworks' distribution of Cloudbreak, please visit <https://hortonworks.com/open-source/cloudbreak/>. For more information on Hortonworks support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Hortonworks reserves the right to change any software described herein at any time, and without notice. Hortonworks assumes no responsibility or liability arising from the use of software described herein, except as expressly agreed to in writing by Hortonworks.

Trademark: Cloudbreak is/are trademark/s of The Apache Software Foundation 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 HORTONWORKS, HORTONWORKS DOES NOT MAKE OR GIVE ANY REPRESENTATION, WARRANTY, OR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH Cloudbreak OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. HORTONWORKS DOES NOT WARRANT THAT Cloudbreak WILL OPERATE UNINTERRUPTED OR THAT IT WILL BE FREE FROM DEFECTS OR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION OR UNAVAILABILITY, OR THAT Cloudbreak WILL MEET ALL OF CUSTOMER'S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, HORTONWORKS 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.