Cloudbreak Release Notes

Date of Publish: 2021-03-16

Contents

About this release

Cloudbreak 2.9.2 is a general availability release: It is suitable for production.

Attention: As of December 31, 2021, Cloudbreak reached end of support. For more information, see Support lifecycle policy. Cloudera recommends that you migrate your workloads to CDP Public Cloud.

Note:

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

New features

Cloudbreak 2.9.2 is a maintenance release with no new features but containing various stability and usability improvements, mainly enhancing the data access layer and the autoscaling component.

There are no new features.

Attention: As of December 31, 2021, Cloudbreak reached end of support. For more information, see Support lifecycle policy. Cloudera recommends that you migrate your workloads to CDP Public Cloud.

Technical preview features

Cloudbreak 2.9.2 includes the following technical preview features. These features should only be used for evaluation and are not suitable for production:

Attention: As of December 31, 2021, Cloudbreak reached end of support. For more information, see Support lifecycle policy. Cloudera recommends that you migrate your workloads to CDP Public Cloud.

Feature	Documentation link
Data lakes	Working with data lakes (TP)
Azure ADLS Gen2	Configuring access to ADLS Gen2
AWS GovCloud	Deploying on AWS GovCloud

Behavioral changes

Cloudbreak 2.9.2 introduces the following changes in behavior as compared to previous Cloudbreak versions.

Attention: As of December 31, 2021, Cloudbreak reached end of support. For more information, see Support lifecycle policy. Cloudera recommends that you migrate your workloads to CDP Public Cloud.

Migrate existing clusters to access new packages behind the paywall

HDP, HDF, and Ambari repositories have been moved behind a paywall. You might have noticed that upscales for existing clusters don't work and no new clusters can be launched from prewarmed images. Refer to the following guidelines, which will help you update the repository settings in order to be able to access the new repositories behind the paywall. See Migrate existing clusters to access new packages behind the paywall.

No default pre-warmed images

This release does not come with a set of pre-warmed images by default, but each supported stack type can be installed by using a base image. Pre-warmed image feature is still supported if you are using either custom images with all software pre-installed and/or are using a local package repository for Ambari and HDP/HDF.

Fixed issues

The following issues were fixed in Cloudbreak 2.9.2:

N/A

⚠️ **Attention:** As of December 31, 2021, Cloudbreak reached end of support. For more information, see Support lifecycle policy. Cloudera recommends that you migrate your workloads to CDP Public Cloud.

Known issues

Cloudbreak 2.9.2 includes the following known issues:

⚠️ **Attention:** As of December 31, 2021, Cloudbreak reached end of support. For more information, see Support lifecycle policy. Cloudera recommends that you migrate your workloads to CDP Public Cloud.

Known issues: Cloudbreak

Issue	Description	Workaround

BUG-117004	When defining network security rules during cluster creation on Azure, when ICMP protocol is used, cluster creation fails with an error similar to:
	
	`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`
		When defining network security rules during cluster creation on Azure, do not use the ICMP protocol.

Issue	Description	Workaround
BUG-117005	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: ``` 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]. ``` 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.	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.
BUG-110998	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.	Click on "Do not configure".
BUG-99581	The Event History in the Cloudbreak web UI displays the following message: Manual recovery is needed for the following failed nodes: [] 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.	If all services are green and healthy in Ambari web UI, then syncing the cluster should fix the problem.
BUG-110999	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.	Your OpenStack admin must import these images manually by using the instructions in Import HDP and HDF images to OpenStack.
BUG-112787	When a cluster with the same name as specified in CLI JSON already exists, CLI returns: ``` ERROR: status code: 403, message: Access is denied. ```	To avoid this error, pass the cluster name as a parameter with cb cluster create instead of including cluster name in the CLI JOSN definition.

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	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 nifi.web.proxy.host blueprint parameter does not get updated to include the new nodes, and as a result the NiFi UI is not reachable from these nodes.	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 nifi.web.proxy.host parameter does not get updated to include the new hosts, and as a result the NiFi UI is not reachable from these hosts. HOST1-IP:PORT,HOST2-IP:PORT,HOST3-IP:PORT

Known issues: Data lake

Issue	Description	Workaround
BUG-109369	Hive does not start on a HDP 2.6 data lake when Kerberos is enabled.	1. Modify /etc/hadoop/<Ambari-version>/0/core-site.xml and /etc/hadoop/conf.backup/core-site.xml by adding the following: ``` <configuration> <property> <name>hadoop.security.authentic name> <value>kerberos</ value> </property> </configuration> ``` 2. Restart affected services.
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: ``` ERROR client.ServiceClient: Error on destroy 'llap0': not found. Failed: org.apache.hadoop.security.AccessControlException: /user/hive/.yarn/ package/LLAP (is not a directory) ```	1. 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. 2. Start HiveServer2.

Image catalog updates

This section lists image catalog updates published for Cloudbreak 2.9.2.

⚠ **Attention:** As of December 31, 2021, Cloudbreak reached end of support. For more information, see Support lifecycle policy. Cloudera recommends that you migrate your workloads to CDP Public Cloud.

Note:

This release does not come with a set of pre-warmed images by default, but each supported stack type can be installed by using a base image. Pre-warmed image feature is still supported if you are using either custom images with all software pre-installed and/or are using a local package repository for Ambari and HDP/HDF.

Default HDP, HDF, and Ambari versions provided with Cloudbreak 2.9.2:

Platform version	Stack version	Ambari version
HDP 2.6	HDP 2.6.5.0-292	Ambari 2.6.2.59
HDP 3.1	HDP 3.1.0.0-78	Ambari 2.7.3.39
HDP 3.1.5	HDP 3.1.5.0-152	Ambari 2.7.5.0
HDF 3.3	HDF 3.3.1.0-10	Ambari 2.7.3.39
HDF 3.5	HDF 3.5.2.0-99	Ambari 2.7.5.0

Legal information

Cloudbreak 2.9.2

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.