

Cloudbreak release notes 2

Cloudbreak Release Notes

Date of Publish: 2019-02-06



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

Contents

About this release.....	3
New features.....	3
Technical preview features.....	4
Behavioral changes.....	4
Fixed issues.....	5
Known issues.....	7
Image catalog updates.....	10
Legal information.....	11

About this release

Cloudbreak 2.9.0 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 2.9.0 introduces the following 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](#).

Feature	Description	Documentation link
Specifying resource group name on Azure	When creating a cluster on Azure, you can specify the name for the new resource group where the cluster will be deployed.	Resource group name
Multiple existing security groups on AWS	When creating a cluster on AWS, you can select multiple existing security groups. This option is available when an existing VPC is selected.	Create a cluster on AWS
EBS volume encryption on AWS	You can optionally configure encryption for EBS volumes attached to cluster instances running on EC2. Default or customer-managed encryption keys can be used.	EBS encryption on AWS
Shared VPCs on GCP	When creating a cluster on Google Cloud, you can place it in an existing shared VPC.	Shared networks on GCP
GCP volume encryption	By default, Google Compute Engine encrypts data at rest stored on disks. You can optionally configure encryption for the encryption keys used for disk encryption. Customer-supplied (CSEK) or customer-managed (CMEK) encryption keys can be used.	Disk encryption on GCP
Workspaces	Cloudbreak introduces a new authorization model, which allows resource sharing via workspaces. In addition to using their personal workspaces, users can create shared workspaces to share resources.	Workspaces
Operations audit logging	Cloudbreak records an audit trail of the actions performed by Cloudbreak users as well as those performed by the Cloudbreak application.	Operations audit logging
Updating long-running clusters	Cloudbreak supports updating base image's operating system and any third party packages that have been installed as well as upgrading Ambari, HDP and HDF.	Updating OS and tools on long-running clusters Updating Ambari and HDP/HDF on long-running clusters
HDP 3.1	Cloudbreak 2.9 introduces default HDP 3.1 blueprints and allows you to create your custom HDP 3.1 blueprints.	Default cluster configurations

Feature	Description	Documentation link
HDF 3.3	Cloudbreak 2.9 introduces default HDF 3.3 blueprints and allows you to create your custom HDF 3.3 blueprints.	Default cluster configurations
Recipe parameters	Supported parameters can be specified in recipes as variables by using mustache kind of templating with "{{{ }}}" syntax.	Writing recipes Recipe parameters
Shebang in Python recipes	Cloudbreak supports using shebang in Python scripts run as recipes.	Writing recipes
AWS GovCloud (TP)	You can install Cloudbreak and create Cloudbreak-managed clusters on AWS GovCloud. This feature is technical preview.	Deploying on AWS GovCloud
Azure ADLS Gen2 (TP)	When creating a cluster on Azure, you can optionally configure access to ADLS Gen2. This feature is technical preview.	Configuring access to ADLS Gen2
New and changed data lake blueprints (TP)	<p>Cloudbreak 2.9 includes three data lake blueprints:</p> <ul style="list-style-type: none"> HDP 2.6 data lake HA blueprint HDP 2.6 data lake blueprint including Atlas HDP 3.1 data lake blueprint <p>The data lake feature remains technical preview.</p> <p> Note: Hive Metastore has been removed from the HDP 3.x data lake blueprints, but setting up an external database allows all clusters attached to a data lake to connect to the same Hive Metastore.</p>	Working with data lakes

Technical preview features

Cloudbreak 2.9.0 introduces 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 2.9.0 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](#).

Affected item	Description
Image catalog	The image catalog was extended to include package versions. Image catalogs which do not include package versions remain valid.
Create cluster wizard > Hardware and Storage	The edit view of the "Hardware and Storage" page in the create cluster wizard was divided into sub-sections to improve readability.
Create cluster wizard > Network	The "Network" page in the create cluster wizard was separated into two pages: "Network and Availability" (Allows you to select network/subnet and specify subnet CIDR) and "Network Security" (Allows you to define security groups).
Create cluster wizard > Recipes	The recipes selection UI was changed: To select an existing recipe, select it from the dropdown and then click "Attach".
History page	The History page was removed and the "History" option is no longer available from the navigation menu.
Navigation menu	The option to collapse the navigation menu was removed.
Data lake	Hive Metastore has been removed from the HDP 3.x data lake blueprints, but setting up an external database allows all clusters attached to a data lake to connect to the same Hive Metastore. There is no change for HDP 2.x data lake blueprints.
Recipes	Unused post-cluster scripts have been cleaned up.

Fixed issues

The following issues were fixed in Cloudbreak 2.9.0:



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](#).

Issue	Description	Category
BUG-85653	Add missing documentation for how to set custom domain and custom hostname when creating a cluster.	Documentation
BUG-110397	The Clusters dashboard is very slow when there are more than 50 items.	Usability
BUG-116758	External database cannot be added for a cluster attached to a data lake.	Usability
BUG-117245	Decrease UI container image sizes by removing cached data.	Usability
BUG-116622	In case of advanced kerberos config with <code>manage_krb5_conf=true</code> cb must not set up a <code>krb5.conf</code> by salt.	Stability
BUG-116582	Cbd LDAP mapping command warning when a remote database is used.	Usability
BUG-116581	Cbd create-bundle command should remove LDAP password from <code>uaa.yml</code> .	Security
BUG-112581	Cloudbreak LDAP configuration overrides "ldap-group-name-attribute" value with "ldap-group-member-attribute".	Stability

Issue	Description	Category
BUG-109889	Cluster create command with HA blueprint fails with "status code: 500, message: Incorrect number of 'HIVE_SERVER' components are in '[master_nn1, master_nn2]' hostgroups: count: 2, min: 1 max: 1" error	Stability
BUG-117357	SLES can't install mpack, missing xmllint.	Stability
BUG-99322	"@#\$%:&*;" in blueprint name causes "duplicate key value violates unique constraint" error.	Usability
BUG-110657	Remove num_llap_nodes configuration from LLAP blueprint. The number of LLAP nodes should be determined by Ambari based on the available memory.	Usability
BUG-105326	GCP quickstart throws resource warning for source_image version.	Usability
BUG-105308	In some cases, when a cluster is terminated while its creation process is still pending, the cluster termination fails with the following exception: "Unable to find com.sequenceiq.cloudbreak.domain.Constraint with id 2250".	Stability
BUG-110505	Blueprint file upload fails the second time if using the same file name.	Usability
BUG-106665	"Region doesn't contain availability zone" exception when availability-zone is not set on AWS.	Usability
BUG-105439	When the master node goes down or is removed, the cluster remains in available status.	Usability
BUG-105309	In some cases, when a cluster is terminated while its creation process is still pending, blueprint remains attached to the cluster, causing the following error when one tries to delete the blueprint: "There is a cluster 'perftest-d0e5q5y7lj' which uses blueprint 'multinode-hdfs-yarn-ez0plywf2c'. Please remove this cluster before deleting the blueprint".	Usability
BUG-110336	Swagger response AmbariRepoDetails has extra mandatory field.	Usability
BUG-105312	When you delete an existing scaling policy or an existing alert, you will see the following confirmation message: "Scaling policy / Alert has been saved". This message should state: "Scaling policy / Alert has been deleted". The deletion occurs correctly, but the confirmation message is incorrect.	Usability
BUG-108976	OS user does not get propagated correctly in Azure quickstart. As a result, the quickstart template times out due to incorrect user.	Stability
BUG-110090	It is possible to delete a data lake cluster while a cluster attached to this data lake exists.	Stability
BUG-106868	No base image is available and therefore no cluster can be created when custom HDF version is larger than 3.1.	Usability

Issue	Description	Category
BUG-109037	Active Directory password that includes the ">" character is not escaped in Knox SSO config.	Stability
BUG-109038	Image base AMI is invalid for centos7.	Stability
BUG-108767	Default post cluster install recipes do not run.	Stability
BUG-108064	Recipe error shown in the UI even when something other than a recipe failed.	Usability
BUG-107643	HDP repo parameters in the UI are incorrect and confusing.	Usability
BUG-106440	Unable to register a database in Cloudbreak when the URL contains a hyphen in the database's name.	Usability
BUG-105997	m5 instance type for AWS should allow to use ST1 (Throughput Optimized HDD) storage.	Usability
BUG-106140	Wrong JSON format is generated if SSO is not configured in Gateway.	Stability
BUG-105498	Cloudbreak UI generates incorrect OpenStack resource links.	Usability
BUG-105205	When the ambari server host was removed from a HA cluster with autorepair on Ambari server hostgroup, no autorepair occurred.	Stability
BUG-105065	Error message is missing in the CLI when ldap server has no port.	Usability
BUG-105445	Document to remove VM dependencies on Azure.	Documentation
BUG-109668	Document minimal Azure role.	Documentation
RMP-11561	Document prerequisites for supported browsers.	Documentation
BUG-105503	Document adding a green SSL certificate to access Ambari.	Documentation

Known issues

Cloudbreak 2.9.0 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-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 2.9.0.

Issue	Description	Workaround
BUG-116919	<p>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:</p> <pre data-bbox="656 348 1166 611">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>	<p>When defining network security group rules on Google Cloud, make sure to define a valid range.</p>
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="656 785 1084 1377">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>

Issue	Description	Workaround
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="657 346 1044 632">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>
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 Import HDP and HDF images to OpenStack.</p>
BUG-112787	<p>When a cluster with the same name as specified in CLI JSON already exists, CLI returns:</p> <pre data-bbox="657 1543 1044 1648">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	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.	<ol style="list-style-type: none"> Modify /etc/hadoop/<Ambari-version>/core-site.xml and /etc/hadoop/conf.backup/core-site.xml by adding the following: <pre><configuration> <property> <name>hadoop.security.authentication.name< <value>kerberos</value> </property> </configuration></pre> 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: <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"> 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. Start HiveServer2.

Image catalog updates

This section lists image catalog updates published for Cloudbreak 2.9.0.



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](#).

Default HDP, HDF, and Ambari versions provided with Cloudbreak 2.9.0:

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 2.9.0

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.