

Cloudbreak release notes 2

Cloudbreak Release Notes

Date of Publish: 2018-09-14



<http://docs.hortonworks.com>

Contents

About this release.....	3
New features.....	3
Behavioral changes.....	3
Fixed issues.....	4
Known issues.....	6
Image catalog updates.....	8
Legal information.....	8

About this release

Cloudbreak 2.8.0 TP is a technical preview release: It is not suitable for production.

New features

Cloudbreak 2.8.0 TP introduces the following new features:

Feature	Description	Documentation
AWS GovCloud	Cloudbreak supports installing Cloudbreak and creating Cloudbreak-managed clusters on AWS GovCloud.	Deploying on AWS vs AWS GovCloud
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
GCP volume encryption	By default, 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
User authorization	Cloudbreak introduces a new authorization model, which allows resource sharing via organizations.	User authorization
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 cluster	Cloudbreak supports updating base image's operating system and any third party packages that have been installed.	Updating long-running clusters
Data lake HA and Atlas support	Cloudbreak includes two data lake blueprints: <ul style="list-style-type: none"> Data lake HA blueprint Data lake blueprint including Atlas (HA is not supported) 	Working with Data Lakes (TP)
Multiple existing security groups on AWS	Multiple existing security groups can be specified when creating a cluster via CLI on AWS.	Multiple existing security groups on AWS
Shebang in Python recipes	Cloudbreak supports using shebang in Python scripts run as recipes.	Writing recipes
HDF 3.2	Cloudbreak can be used to deploy HDF 3.2 clusters by using one of the two default HDF 3.2 blueprints: <ul style="list-style-type: none"> Flow Management clusters with Apache NiFi Messaging clusters with Apache Kafka. 	Default cluster configurations

Behavioral changes

Cloudbreak 2.8.0 TP introduces the following changes in behavior as compared to previous Cloudbreak versions:

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 in the navigation menu.
Navigation menu	The option to collapse the navigation menu was removed.

Fixed issues

The following issues were fixed in Cloudbreak 2.8.0 TP:

Issue	Description	Category	Fix version
BUG-99322	"@#% :&*;" in blueprint name causes "duplicate key value violates unique constraint" error.	Usability	2.8.0
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	2.8.0
BUG-105326	GCP quickstart throws resource warning for source_image version.	Usability	2.8.0
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 Constraint	2.8.0
BUG-110505	Blueprint file upload fails the second time if using the same file name.	Usability	2.8.0
BUG-106665	"Region doesn't contain availability zone" exception when availability-zone is not set on AWS.	Usability	2.8.0
BUG-105439	When the master node goes down or is removed, the cluster remains in available status.	Usability	2.8.0

Issue	Description	Category	Fix version
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	2.8.0
BUG-110336	Swagger response AmbariRepoDetails has extra mandatory field.	Usability	2.8.0
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	2.8.0
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	2.8.0
BUG-110090	It is possible to delete a data lake cluster while a cluster attached to this data lake exists.	Stability	2.8.0
BUG-106868	No base image is available and therefore no cluster can be created when custom HDF version is larger than 3.1.	Usability	2.8.0
BUG-109037	Active Directory password that includes the ">" character is not escaped in Knox SSO config.	Stability	2.8.0
BUG-109038	Image base AMI is invalid for centos7.	Stability	2.8.0
BUG-108767	Default post cluster install recipes do not run.	Stability	2.8.0
BUG-108064	Recipe error shown in the UI even when something other than a recipe failed.	Usability	2.8.0
BUG-107643	HDP repo parameters in the UI are incorrect and confusing.	Usability	2.8.0
BUG-106440	Unable to register a database in Cloudbreak when the URL contains a hyphen in the database's name.	Usability	2.8.0
BUG-105997	m5 instance type for AWS should allow to use ST1 (Throughput Optimized HDD) storage.	Usability	2.8.0
BUG-106140	Wrong JSON format is generated if SSO is not configured in Gateway.	Stability	2.8.0

Issue	Description	Category	Fix version
BUG-105498	Cloudbreak UI generates incorrect OpenStack resource links.	Usability	2.8.0
BUG-105205	When the ambari server host was removed from a HA cluster with autorepair on Ambari server hostgroup, no autorepair occurred.	Stability	2.8.0
BUG-105065	Error message is missing in the CLI when ldap server has no port.	Usability	2.8.0
BUG-105445	Document to remove VM dependencies on Azure.	Documentation	2.8.0
BUG-109668	Document minimal Azure role.	Documentation	2.8.0
RMP-11561	Document prerequisites for supported browsers.	Documentation	2.8.0
BUG-105503	Document adding a green SSL certificate to access Ambari.	Documentation	2.8.0

Known issues

Cloudbreak 2.8.0 TP includes the following known issues:

Known issues: Cloudbreak

Issue	Description	Workaround
RMP-11665	Due to Ambari 2.7.0.0 lack of support for Amazon Linux 2, it is not possible to launch HDF 3.2 clusters that use Ambari 2.7.0.0 on AWS.	There is no workaround, but you can: <ul style="list-style-type: none"> • Use HDF 3.1 instead HDF 3.2. • Use Ambari 2.7.1 (once released) with HDF 3.2. • Use HDF 3.2 with a different cloud provider.
BUG-110998	When creating a workload 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 not be available for any clusters other than data lakes.	Click on "Do not configure".
BUG-109369	Hive does not start on a data lake when Kerberos is enabled.	<ol style="list-style-type: none"> 1. Modify /etc/hadoop/<VERSION>/0/core-site.xml and /etc/hadoop/conf.backup/core-site.xml by adding the following: <pre><configuration> <property> <name>hadoop.security.authentication.name</name> <value>kerberos</value> </property> </configuration></pre> 2. Restart affected services.

Issue	Description	Workaround
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-110397	The Clusters dashboard is very slow when there are more than 50 items.	
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 .
N/A	The following Azure and GCP advanced cluster options are missing from the web UI: Don't create public IP, Don't create firewall rules.	These web UI options are available in Cloudbreak 2.7.x and they will be reinstated in the next release. In Cloudbreak 2.8, you can only set these by using CLI.

Known issues: Ambari

The known issues described here were discovered when testing Cloudbreak with Ambari versions that are used by default in Cloudbreak. For general Ambari and known issues, refer to Ambari release notes.

Issue	Description	Workaround
BUG-109369	Hive does not start on a data lake when Kerberos is enabled.	<ol style="list-style-type: none"> 1. Modify /etc/hadoop/<Ambari-version>/0/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> 2. Restart affected services.

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.

There are no known issues related to HDP.

Known issues: HDF

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 <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.	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. HOST1-IP:PORT,HOST2-IP:PORT,HOST3-IP:PORT

Image catalog updates

This section lists image catalog updates published for this Cloudbreak version.

Default HDP, HDF, and Ambari versions provided with Cloudbreak 2.8.0 TP:

Default versions used for HDP 2.6:

- HDP 2.6.5.0-292
- Ambari 2.6.2.2

Default versions used for HDF 3.2:

- HDF 3.2.0.0-520
- Ambari 2.7.0.0

Legal information

Cloudbreak 2.8.0 TP

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.