

Data Hub

Troubleshooting Data Hub

Date published: 2019-12-17

Date modified: 2023-06-27

CLOUDERA

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

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks 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 CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA 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.

Contents

Checking logs.....	4
Cloudera Manager logs.....	4
Recipe logs.....	4
 Troubleshooting cluster creation.....	 4
Failed to retrieve the server's certificate.....	4
Quota limitations.....	4
 Troubleshooting upgrade issues.....	 5

Checking logs

When troubleshooting Data Hub clusters, you can access the following logs:

Cloudera Manager logs

Data Hub uses Cloudera Manager to orchestrate the installation of Cloudera Runtime components. Each instance in the cluster runs an Cloudera Manager agent which connects to the Cloudera Manager server. Cloudera Manager server is declared by the user during the cluster installation wizard.

Refer to the following documentation for information on how to access Cloudera Manager logs:

Related Information

[Cloudera Manager Logs](#)

Recipe logs

Data Hub supports “recipes” - user-provided customization scripts that can be run prior to or after cluster installation. It is the user’s responsibility to provide an idempotent well tested script. If the execution fails, the recipe logs can be found at `/var/log/recipes` on the nodes on which the recipes were executed.

It is advised, but not required to have an advanced logging mechanism in the script, as Data Hub always logs every script that are run. Recipes are often the sources of installation failures as users might try to remove necessary packages or reconfigure services.

Troubleshooting cluster creation

The following section lists common issues related to cluster creation.

Failed to retrieve the server's certificate

Data Hub cluster fails with "Infrastructure creation failed. Reason: Failed to retrieve the server's certificate".

When a cluster fails with "Infrastructure creation failed. Reason: Failed to retrieve the server's certificate", the most common reasons for this error are related to using your own custom image. If you are using your own custom image and if using the CLI, you must send the `imageId` explicitly in the CLI in the cluster template request.

Quota limitations

Each cloud provider has quota limitations on various cloud resources, and these quotas can usually be increased on request. If there is an error message in Data Hub stating that there are no more available EIPs (Elastic IP Address) or VPCs, you need to request more of these resources.

To see the limitations visit the cloud provider’s site:

- [AWS service limits](#)
- [Azure subscription and service limits, quotas, and constraints](#)

Troubleshooting upgrade issues

The following section lists common issues related to upgrading Data Hubs.

Failed to start role of service profiler_scheduler in cluster

While upgrading from Cloudera Runtime version 7.2.17 to 7.2.18, and specifically during the Data Hub OS upgrade step, the cluster goes into the failure state. The following message is seen:

```
NODE_FAILURE:
New node(s) could not be added to the cluster. Reason Please find more details on Cloudera Manager UI. Failed command(s): Start(id=1546339088): Failed to start role profc6cf3856-PROFILER_SCHEDULER_AGENT-484032cb8f17cacf9e684efe50 of service profiler_scheduler in cluster cdp-dc-profilers-258395ef.
```

Impact on Data Catalog profilers:

If the Data Hub is not created, the profilers will not be created in Cloudera Runtime 7.2.18 version.

Workaround:

You must use the following process to bring up the Data Catalog profilers in the 7.2.18 version.

1. First you must delete your existing 7.2.17 clusters. For more information, see [Deleting profiler cluster](#).
2. Next, after you upgrade to the 7.2.18 Data Lake, then you can launch the Data Catalog profilers. For more information, see [Launch profiler cluster](#).



Note: There is no data loss expected on the users' side or the Profiler analysis. However, the only loss that could be expected is related to the last runtime value of the profiler and the profiler run history. The Profiler Last Runtime history refers to the records of how many runs of the profiler are displayed on the history page. It includes information on whether the runs were completed successfully or resulted in failures.