

Troubleshooting Cloudera Machine Learning

Date published: 2020-07-16

Date modified: 2024-11-21



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

- Troubleshooting.....4**
- Preflight checks.....4**
 - Instance type preflight check fails..... 14
- Cloudera Machine Learning service with Data Lake upgrades.....15**
- Common Cloudera Machine Learning errors and solutions..... 16**
- Troubleshooting Cloudera Machine Learning Workspaces on AWS..... 19**
- Troubleshooting Cloudera Machine Learning Workspaces on Azure..... 21**
- Logs for Cloudera Machine Learning Workspaces..... 21**
 - Downloading diagnostic bundles for a workspace..... 22
- Troubleshooting issues with workloads.....23**
 - Troubleshooting Spark issues.....23
 - Troubleshooting Kerberos issues.....24

Troubleshooting

This topic describes a recommended series of steps to help you start diagnosing issues with a Cloudera Machine Learning Workspace.

- **Issues with Provisioning Cloudera Machine Learning Workspaces:** If provisioning an Cloudera Machine Learning Workspace fails, first go to your cloud provider account and make sure that you have all the resources required to provision an Cloudera Machine Learning Workspace. If failures persist, start debugging by reviewing the error messages on the screen. Check the workspace logs to find where the problem originates from. For more details on the troubleshooting resources available to you, see [Troubleshooting Cloudera Machine Learning Workspaces on AWS](#) on page 19.
- **Issues with Accessing Cloudera Machine Learning Workspaces:** If your Cloudera Machine Learning Admin has already provisioned a workspace for you but attempting to access the workspace fails, confirm with your Cloudera Machine Learning Admin that they have completed all the steps required to grant you access. See: [Configuring User Access to Cloudera Machine Learning](#)
- **Issues with Running Workloads:** If you have access to a workspace but are having trouble running sessions/jobs/experiments, and so on, see if your error is already listed here: [Troubleshooting issues with workloads](#) on page 23.

Cloudera Support

If you need assistance, contact Cloudera Support. Cloudera customers can register for an account to create a support ticket at the [support portal](#). For Cloudera Data Platform issues in particular, make sure you include the Request ID associated with your error message in the support case you create.

Preflight checks

Part of provisioning and managing a Kubernetes cluster in the cloud is ensuring that your account and your environment is properly configured before beginning. The Cloudera Data Platform automatically runs a series of preflight checks which can help in determining if there is a problem before creating new resources or adjusting existing ones.

When creating a brand new cluster, preflight validation checks are used to ensure that the resources you are requesting as well as your environment are ready and configured correctly. For existing clusters, many of the infrastructure validations are skipped since they are not needed anymore. Instead, the validations which concern the resources being adjusted are executed.

Results

The result of running a collection of preflight checks is represented as a single aggregated value that lets you know whether it is safe to proceed with your actions. Each individual validation which was run is included in the response.

A preflight validation contains information concerning what it was checking for and what the result of that check was. For example, this is a very simple check to ensure that the type of node pool image is valid in the region:

Name	Instance type
Description	Instance groups must have an instance type that exists in the region in which they will be created. For EKS, there is additional verification for EKS support and usage class.
Category	COMMON

Name	Instance type
Status	FAILED
Message	The instance type validation failed.
Detailed Message	Instance type validation failed for Standard_B2s1 in west us2. Check to ensure that this instance type is valid in that region.
Duration	659 ms

PASSED

The validation successfully passed all criteria.

WARNING

The validation was either unable to fully check all of its criteria or it found a potential issue which could affect the success of the operation. This type of failure will not stop a cluster from being created or modified, but it does require further investigation.

SKIPPED

The validation was skipped because it does not apply to the current request. This can happen for many reasons, such as using a cloud provider that is not applicable to the preflight check.

FAILED

The validation failed its expected criteria and provisioning or updating of a cluster cannot proceed. In this case, there is an identified problem that needs resolution before continuing.

List of preflight checks

Node / Agent Pools

The following preflight validation checks pertain to the node / agent instance groups which are created to launch new nodes inside of the Kubernetes cluster.

Name	Description	Cloud	Type
Instance Count	The number of nodes in each instance group must not exceed a predefined threshold. Remediation: Change the requested size of an instance group to be within the boundaries specified by the error message.	All	Create Update
Group Count	The number of distinct node/agent pool groups must not exceed a predefined threshold. Remediation: Change the number of distinct instance groups to be within the	All	Create Update

Name	Description	Cloud	Type
	boundaries specified by the error message.		
Instance Naming	<p>The name of each instance group is restricted by cloud providers. There are differences in the size and characters allowed by each provider.</p> <p>Remediation: Change the name of the instance group to conform to the cloud provider's requirements. This could include changing the overall length of the name or only using certain approved characters.</p>	All	Create Update
Instance Type	<p>Instance image types are not universal across all regions. Some providers further restrict this to service level, and whether they can be used for Kubernetes.</p> <p>Remediation: Choose a different region or service type for the desired instance type. If this is not possible, then a different instance type must be chosen.</p>	All	Create Update
Kubernetes Version	<p>Each cloud provider supports different versions of Kubernetes. Cloudera Data Platform has also only been certified to work with particular versions.</p> <p>Remediation: Choose a different version of Kubernetes that is supported by your cloud provider in the region you are deploying in.</p> <p>https://docs.aws.amazon.com/eks/latest/userguide/kubernetes-versions.html</p>	All	Create

Name	Description	Cloud	Type
Placement Rules	<p>Some instance types are not allowed to be grouped within a single availability zone.</p> <p>Remediation: Remove the restriction on single availability zone, or choose a different instance type.</p>	AWS	Create Update

Infrastructure

The following preflight validation checks pertain to Cloudera's control plane infrastructure and your specific account within that control plane.

Name	Description	Cloud	Type
Restricted IAM Policies	<p>If your account is trying to provision with restricted IAM policies, then it needs to have those policies defined before deploying the cluster.</p> <p>Remediation: Check Cloudera's documentation on restricted IAM policies to ensure that you have the correctly named policies defined and are accessible.</p>	AWS	Create
Proxy Connectivity	<p>When provisioning a private cluster, your environment must have the cluster proxy enabled and it must be healthy.</p> <p>Remediation: Create a new environment which has the cluster proxy service enabled or check that your existing FreeIPA Virtual Machine is running and healthy.</p> <p>https://docs.cloudera.com/management-console/test/connection-to-private-subnets/topics/mc-ccm-overview.html</p>	All	Create Update
Data Lake Connectivity	A healthy data lake with a functioning FreeIPA DNS server is required in order to provision a new cluster.	All	Create

Name	Description	Cloud	Type
	<p>Remediation: Check to ensure that the FreeIPA DNS server is running and healthy inside of your network.</p> <p>https://docs.cloudera.com/management-console/cloud/data-lakes/topics/mc-data-lake.html</p> <p>https://docs.cloudera.com/management-console/cloud/identity-management/topics/mc-identity-management.html</p>		

Networking

The following preflight validation checks pertain to specific network configurations inside of the cloud provider.

Name	Description	Cloud	Type
Shared Virtual Private Cloud	<p>When a Virtual Private Cloud is shared between multiple subscriptions, access to modify this Virtual Private Cloud needs to be granted.</p> <p>Remediation: Check the permission for which roles can make modifications to the Virtual Private Cloud</p> <p>https://docs.aws.amazon.com/vpc/latest/userguide/vpc-sharing.html</p> <p>https://docs.cloudera.com/cdp-public-cloud/cloud/requirements-aws/topics/mc-aws-req-vpc.html</p>	AWS	Create
Subnet Availability Zone	<p>All subnets which are part of the environment must be located in at least 2 different Availability Zones.</p> <p>Remediation: Recreate your environment and choose subnets that satisfy the requirement of being in at least 2 different Availability Zones.</p>	AWS	Create

Name	Description	Cloud	Type
Subnet Load Balancer Tagging	<p>In order for load balancers to choose subnets correctly, a subnet needs to have either the public or private ELB tags defined.</p> <p>Remediation: Tag subnets with either <code>kubernetes.io/role/elb</code> or <code>kubernetes.io/role/internal-elb</code> based on whether they are public or private.</p> <p>https://docs.cloudera.com/cdp-public-cloud/cloud/requirements-aws/topics/mc-aws-req-vpc.html</p>	AWS	Create
API Server Access	<p>Validates that there are no conflicting requests between Kubernetes API CIDR ranges and private AKS clusters.</p> <p>Remediation: When using a private AKS cluster, Kubernetes API CIDR ranges are not supported,</p>	All	Create Update
Available Subnets	<p>Subnets cannot be shared when provisioning Kubernetes clusters on Azure. At least one available subnet must exist that is not being used by another AKS cluster and must not have an existing route table with conflicting pod CIDRs.</p> <p>Remediation: Create a new subnet to satisfy this requirement or delete an old and unused cluster to free an existing subnet.</p>	Azure	Create
Delegated Subnet	<p>A subnet which has been delegated for a particular service cannot be used to provision an Azure AKS cluster.</p> <p>Remediation: Choose a different subnet or remove the delegated service from</p>	All	Create

Name	Description	Cloud	Type
	at least one subnet in the environment.		
Kubernetes API Server Security	Validates that the supplied IP CIDR ranges are valid and do not overlap any reserved IP ranges. Each cloud provider has a limit set on the maximum number of allowed CIDRs. Remediation: Use valid CIDR formats and ranges when limiting access to the Kubernetes API server and limit the number of ranges specified.	All	Create
Kubernetes Service CIDR Validation	Validates that the specified service CIDR for Kubernetes services does not overlap any restricted CIDR ranges and is a valid CIDR format. Remediation: Change the service CIDR so that it does not conflict with any pod CIDRs or other routes on the subnet.	All	Create
Autoscale Parameters	Azure's built-in autoscaler has limitations on the ranges of values for scale-up and scale-down operations. Remediation: Adjust the specified parameters from the error message which are not within the required ranges.	Azure	Create Update

Example

```
{
  "result": "PASSED",
  "summary": {
    "passed": 8,
    "warning": 0,
    "failed": 0,
    "skipped": 10,
    "total": 18
  },
  "message": "The cluster validation has passed, but some checks were skipped",
  "validations": [
```

```

    {
      "name": "Instance Count",
      "description": "Each instance count must be between minInstance and maxInstance inclusively. The minInstance and maxInstance of infrastructure group should comply with minimum number of infra nodes and maximum number of infra nodes.",
      "category": "COMMON",
      "status": "PASSED",
      "message": "The minimum and maximum instance counts are correct for all instance groups.",
      "detailedMessage": "The minimum and maximum instance counts are correct for all instance groups.",
      "duration": "1µs"
    },
    {
      "name": "Instance Group Count",
      "description": "Total instance group count must be less than or equal to maximum instance group limit.",
      "category": "COMMON",
      "status": "PASSED",
      "message": "The number of instance groups in the request is less than or equal to the maximum allowed.",
      "detailedMessage": "The total instance group count of 2 is within the limit.",
      "duration": "3µs"
    },
    {
      "name": "Instance Group Naming",
      "description": "Each instance group name must conform the restrictions of the cloud provider. This includes using valid characters and adhering to length restrictions.",
      "category": "COMMON",
      "status": "PASSED",
      "message": "All instance groups meet the naming restrictions for Azure.",
      "detailedMessage": "All instance groups meet the naming restrictions for Azure.",
      "duration": "12µs"
    },
    {
      "name": "Instance Type",
      "description": "Instance groups must have an instance type that exists in the region in which they will be created. For EKS, there is additional verification for EKS support and usage class.",
      "category": "COMMON",
      "status": "PASSED",
      "message": "All instance groups have valid instance types.",
      "detailedMessage": "The following instance types were validated for westus2: Standard_B2s",
      "duration": "599ms"
    },
    {
      "name": "Kubernetes Version",
      "description": "Each cloud provider (Amazon, Azure, Google, etc) supports different versions of Kubernetes.",
      "category": "COMMON",
      "status": "PASSED",
      "message": "The specified Kubernetes version 1.18 has been resolved to 1.18.17 and is valid on Azure",
      "detailedMessage": "The specified Kubernetes version 1.18 has been resolved to 1.18.17 and is valid on Azure",
      "duration": "388ms"
    }
  ],
  {

```

```

        "name": "Placement Rule",
        "description": "Instance Types must be allowed by the placement r
ule.",
        "category": "COMMON",
        "status": "SKIPPED",
        "message": "Skipping validation since the cloud platform is Azure
.",
        "detailedMessage": "Skipping validation since the cloud platform
is Azure.",
        "duration": "6µs"
    },
    {
        "name": "Entitlement Check",
        "description": "When the entitlement LIFTIE_USE_PRECREATED_IAM_RE
SOURCES is enabled, the expected profile (cdp-liftie-instance-profile) shoul
d exist and it needs to have the necessary roles attached to it. ",
        "category": "ENTITLEMENTS",
        "status": "SKIPPED",
        "message": "IAM Resource Entitlement validation skipped for Cl
oud Provider azure.",
        "detailedMessage": "IAM Resource Entitlement validation skipped
for Cloud Provider azure.",
        "duration": "14µs"
    },
    {
        "name": "Cluster Proxy Connectivity",
        "description": "Verifies connectivity to the cluster proxy ser
vice which is used to register private cluster endpoints.",
        "category": "CONTROL_PLANE",
        "status": "SKIPPED",
        "message": "Connectivity to the cluster connectivity manager will
be skipped since this is not a private cluster.",
        "detailedMessage": "The cluster being provisioned is not marked
as private in the provisioning request.",
        "duration": "1µs"
    },
    {
        "name": "Cluster Proxy Enabled",
        "description": "Verifies that the environment was created with th
e cluster proxy service enabled.",
        "category": "CONTROL_PLANE",
        "status": "SKIPPED",
        "message": "Skipping the environment check for cluster proxy c
onnectivity since the cluster is public.",
        "detailedMessage": "Skipping the environment check for cluster
proxy connectivity since the cluster is public.",
        "duration": "1µs"
    },
    {
        "name": "DataLake Connectivity",
        "description": "Validates whether DataLake connection is reach
able and if FreeIPA is available.",
        "category": "CONTROL_PLANE",
        "status": "PASSED",
        "message": "DataLake validation succeeded.",
        "detailedMessage": "Data lake is healthy and reachable. Service
Discovery Feature is enabled, verified DNS entries retrieved for Data Lake
s. Datalake URL : localhost:8081 Service discovery URL : localhost:8082 "
    },
    {
        "name": "AWS Shared VPC Access",
        "description": "When a shared VPC is used, proper access should
be granted.",
        "category": "NETWORK",

```

```

        "status": "SKIPPED",
        "message": "Skipping validation since the cloud platform is Azur
e.",
        "detailedMessage": "Skipping validation since the cloud platform
is Azure.",
        "duration": "3µs"
    },
    {
        "name": "AWS Subnet Availability Zones",
        "description": "When existing AWS subnets are provided for provi
sioning an EKS cluster, the subnets must be in at least 2 different Availabi
lity Zones.",
        "category": "NETWORK",
        "status": "SKIPPED",
        "message": "Skipping validation since the cloud platform is Azure
.",
        "detailedMessage": "Skipping validation since the cloud platform
is Azure.",
        "duration": "2µs"
    },
    {
        "name": "AWS Subnet Tagging",
        "description": "In order for load balancers to choose subnets co
rrectly a subnet needs to have either the public or private ELB tags defined
.",
        "category": "NETWORK",
        "status": "SKIPPED",
        "message": "Skipping validation since the cloud platform is Azu
re.",
        "detailedMessage": "Skipping validation since the cloud platform
is Azure.",
        "duration": "26µs"
    },
    {
        "name": "Azure API Access Parameters",
        "description": "Verifies that the security parameters for locking
down access to the Azure Kubernetes API Server are correct.",
        "category": "NETWORK",
        "status": "PASSED",
        "message": "The API server access parameters specified in the cl
uster request are valid.",
        "detailedMessage": "The cluster will be provisioned as public
with the following whitelist CIDRs: ",
        "duration": "48µs"
    },
    {
        "name": "Azure Available Subnets",
        "description": "When an existing Azure subnet is chosen for pro
visioning an AKS cluster, the subnet must not be in use by any other cluster
. This is a restriction of Kubenet, which is the CNI used on the new cluster
. Although the subnet may have a routing table, it may not have any existing
IP address associations.",
        "category": "NETWORK",
        "status": "PASSED",
        "message": "At least 1 valid subnet was found and can be used for
cluster creation.",
        "detailedMessage": "The cluster can be provisioned using subnet
liftie-dev.internal.2.westus2 in virtual network liftie-dev and resource gr
oup liftie-test",
        "duration": "917ms"
    },
    {
        "name": "Kubernetes API Server CIDR Security",

```

```

      "description": "CIDR blocks for whitelisting access to the Kube
rnetes API Server must not overlap restricted IP ranges.",
      "category": "NETWORK",
      "status": "SKIPPED",
      "message": "Skipping CIDR validation for whitelisting because it
is not enabled.",
      "detailedMessage": "The ability to secure access the Kubernetes
API server via a list of allowed CIDRs is not enabled. This can be enabled
either in the controlplane (currently false) or via the provisioning reques
t (currently false).",
      "duration": "11µs"
    },
    {
      "name": "Service CIDR Validation",
      "description": "CIDR blocks that Kubernetes assigns service IP
addresses from should not overlap with any other networks that are peered or
connected to existing VPC.",
      "category": "NETWORK",
      "status": "SKIPPED",
      "message": "Service CIDR is missing in Network Profile.",
      "detailedMessage": "VPC validation is executed only if the VPC a
lready exists and a service CIDR is specified in the network profile.",
      "duration": "426µs"
    },
    {
      "name": "Azure Autoscale Parameters",
      "description": "The following autoscale parameters for Azure,
which are specified during provisioning and update, need to be in multiples
of 60s. Autoscale parameters: scaleDownDelayAfterAdd, scaleDownDelayAfterFai
lure, scaleDownUnneededTime, scaleDownUnreadyTime.",
      "category": "DEPLOYMENT",
      "status": "SKIPPED",
      "message": "There are no autoscale parameters specified in the
request.",
      "detailedMessage": "The request did not contain an Autoscaler
structure.",
      "duration": "0s"
    }
  ]
}

```

Instance type preflight check fails

On AWS, in a given region, the corresponding Availability Zones (AZs) do not always support the same instance types. The instance type preflight check will fail if the desired instance type is not supported in all AZs in the region.

If the preflight check fails, check all failed AZs for the supported instance types using this command:

```

aws ec2 describe-instance-type-offerings --location-type availability-zone -
--filters Name=location,Values=<az>
--query "InstanceTypeOfferings[ ].InstanceType" --region <region>

```

Remediation:

It is possible, but not guaranteed, that skipping the validation will result in successful provisioning of a Cloudera Machine Learning Workspace. AWS selects the AZ for creating an instance group based on the VPC and Subnet configuration of the customer's datalake; the user cannot choose the AZ.

Alternatively, select an instance type that is supported by all AZs in the region.

Cloudera Machine Learning service with Data Lake upgrades

Cloudera Data Platform environments have two services which can be upgraded individually, the FreeIPA service and the Data Lake (DL) service. Cloudera Machine Learning Workspaces run in Cloudera Data Platform environments. The FreeIPA service provides identity management, and the Data Lake service provides SDX capabilities to Cloudera Machine Learning Workspaces.

In this document we provide FAQs for the behavior of Cloudera Machine Learning Workspaces during a Data Lake upgrade. For information on FreeIPA upgrades, see *Upgrade FreeIPA*.

What kinds of Data Lake upgrades are possible?

The Data Lake service supports the following upgrades.

- Hotfix upgrades
- Version upgrades
- OS version upgrades

These upgrades can all be done from the Cloudera Data Platform Data Lake service UI, or with the Cloudera Data Platform CLI. DL upgrades require downtime. DL upgrades preserve the state of the data lake during the upgrade.

During DL upgrades, the shape of the data lake cannot change. For example, we cannot change a LIGHT_DUTY shape to a MEDIUM_DUTY_HA shape during the DL upgrade.

What is Data Lake migration and is it supported?

To change the DL shape, perform a DL migration, also called DL scaling. For example, a DL migration can change a DL LIGHT_DUTY shape to a MEDIUM_DUTY_HA shape.

[DL migration is currently available as a technical preview](#). Speak to your Cloudera account team to discuss whether it may be suitable for your situation.

What happens if the Data Lake upgrade/migration fails?

There is no automated backup and restore process for the DL. It is recommended to perform a backup of the DL before starting the upgrade process. If the DL upgrade fails, the recommended option is to delete the failed DL using the Cloudera Data Platform CLI (`cdp datalake delete-datalake --datalake-name <dl name>`) and recreate the DL using the `cdpcli`. Once the DL is recreated, you need to restore the DL state from the backup. For more information, refer to *Backup and Restore for the Data Lake*.

Do not delete the environment service during a failed DL upgrade process. Deleting an environment causes all Cloudera Machine Learning Workspace running in this environment to be unusable.

Can the environment service be deleted and recreated at any point if the Data Lake upgrade or migration process has an error?

No, environments with experiences running inside them cannot be deleted at any time. If you delete the environment, then all the experiences (such as a Cloudera Machine Learning Workspace) need to be deleted.

Unless a Cloudera Machine Learning Workspace is first backed up and restored, then all state information is lost and you need to start from a fresh workspace. For more information, see *Backing up Cloudera Machine Learning Workspaces*.

What are the Cloudera Machine Learning Workspace prerequisites for Data Lake upgrades/migrations?

Do the following before upgrading or migrating a Data Lake.

- Upgrade Cloudera Machine Learning Workspaces to the latest version (if an upgrade is available).
- Stop any jobs, sessions, experiments or any workloads that need DL access before performing a DL upgrade.
- Announce to the team that there will be planned downtime for Cloudera Machine Learning Workspaces during the DL upgrade process.

Are the Cloudera Machine Learning Workspaces operational during Data Lake upgrades/migrations?

It is recommended NOT to use Cloudera Machine Learning Workspaces during DL upgrades.

However, the observed behavior of Cloudera Machine Learning Workspaces during DL upgrades is as follows.

- Cloudera Machine Learning Workspaces remain accessible during DL upgrades. Users can log in to a Cloudera Machine Learning Workspace.
- Users can launch sessions, run jobs, experiments, models, and so on which do not require DL access. For example, jobs that do not require IDBroker or SDX/HMS access will function normally.
- Any compute instance that requires IDBroker or SDX/HMS access will fail.
- Any scheduled jobs that require IDBroker or SDX access will fail.
- With regular DL upgrades and migrates, you should suspend Cloudera Machine Learning Workspaces. This require downtime.
- With Zero Downtime Upgrade (ZDU) of DL, Cloudera Machine Learning Workspaces will be functional during upgrades. If the DL version is 7.2.17 or lower, you may encounter the [CDPD-66549](#) known issue.

Are any changes to Cloudera Machine Learning Workspaces needed after a DL is upgraded or migrated successfully?

No further actions are required on a Cloudera Machine Learning Workspace after a successful DL upgrade. Cloudera Machine Learning Workspaces continue to function normally. Make sure to announce to the team that they can start using Cloudera Machine Learning Workspaces as usual.

Related Information

[Backup and restore for the Data Lake](#)

[Backing up ML workspaces](#)

[Upgrade FreeIPA](#)

Common Cloudera Machine Learning errors and solutions

The following sections describe recommended steps to start debugging common error messages you might see in the workspace logs (found under EventsView Logs).

Before you begin

Make sure you have reviewed the list of resources available to you for debugging on Cloudera Machine Learning and AWS: [Troubleshooting Cloudera Machine Learning Workspaces on AWS](#) on page 19

Timezone not properly set for scheduled job

If the timezone is no longer correctly set for a scheduled job, then you should simply set it again. Go to Job Settings, edit the timezone, and update the job.

AWS account resource limits exceeded (Compute, VPC, etc.)

Cloudera Machine Learning Workspace provisioning fails because Cloudera Data Platform could not get access to all the AWS resources needed to deploy a Cloudera Machine Learning Workspace. This is likely because your AWS account either does not have access to those resources or is hitting the resource limits imposed on it.

Sample errors include (from EventsView Logs):

```
Failed to provision cluster. Reason: Failed to wait for provisioner: Wait for status failed with status CREATE_FAILED: error creating eks cluster (cause: InvalidParameterException: Provided subnets subnet-0a648a0cc5976b7a9 Free IPs: 0 , need at least 3 IPs in each subnet to be free for this operation
```

```
Failed to mount storage. Reason: Failed to create mount target: NoFreeAddressesInSubnet: The specified subnet does not have enough free addresses to satisfy the request.
```

AWS accounts have certain hard and soft resource limits imposed on them by default. For example, certain CPU/GPU instances that Cloudera Machine Learning allows you to provision might even have an initial default limit of 0 (set by AWS). This means if you attempt to provision a cluster with those instance types, your request will fail.

Aside from the CPU and GPU compute resource limits, there are other types of limits you can run into. For example, the second error shows that the subnets in your VPC do not have any more free IP addresses for the workspace (and each of the underlying Kubernetes pods). This occurs if the CIDR range mentioned while registering the environment was not large enough for your current needs.

You can use the AWS console to request an increase in limits as needed. Go to the AWS console for the region where the environment was provisioned and then navigate to EC2 Limits.

For networking failures, navigate to EC2VPC. Search for the environment's VPC ID (available on environment Summary page) to see the list of available IP addresses for each subnet. Request more resources as needed.

Related AWS documentation: [AWS Service Limits](#), [Amazon EC2 Resource Limits](#), [EKS Cluster VPC Considerations](#), [AWS CNI Custom Networking](#).

Access denied to AWS credential used for authentication

The cloud credential used to give Cloudera Data Platform access to your AWS account failed authentication. Therefore, Cloudera Data Platform could not provision the resources required to deploy a Cloudera Machine Learning Workspace.

Sample error (from EventsView Logs):

```
Failed to provision storage. Reason: Failed to create new file system: AccessDenied: User: arn:aws:iam::1234567890:user/cross-account-trust-user is not authorized to perform:
```

Your cloud credential gives Cloudera Data Platform access to the region and virtual network that make up the environment thus allowing Cloudera Data Platform to provision resources within that environment. If authentication fails, go to your environment to see how the cloud credentials were set up and confirm whether your account has the permission to perform these actions.

Cloudera Machine Learning installation failures

While the steps to provision resources on AWS were completed successfully, the Cloudera Machine Learning Workspace installation on EKS failed.

Sample error (from EventsView Logs):

```
Failed to install ML workspace. Reason:Error: release mlx-mlx failed: timed out waiting for the condition
```

If you are an advanced user, you can log in to the underlying EKS cluster and use `kubectl` to investigate further into which pods are failing.



Note: This error might be an indication that DNS has been turned off for the VPC. Go to the AWS console for the region where the environment was provisioned and then navigate to EC2 Load Balancers to confirm that DNS is configured properly for the environment's VPC.

Related AWS documentation: [EKS and kubect](#)

Failures due to API throttling

These errors can be harder to prepare for due to their seemingly random nature. Occasionally, AWS will block API calls if it receives too many requests at the same time. For example, this can occur when multiple users are attempting to provision/delete/upgrade clusters at the same time.

Sample error (from EventsView Logs):

```
Failed to delete cluster. Reason: Failed to wait for deletion: Wait for status failed with status DELETE_FAILED: Throttling: Rate exceeded
```

Currently, if you see a 'Throttling: Rate exceeded' error, our recommendation is that you simply try again later.

Related AWS documentation: [AWS API Request Throttling](#)

De-provisioning failures

De-provisioning operations can sometimes fail if AWS resources are not terminated in the right order. This is usually due to timing issues where certain resources might take too long to terminate. This can result in a cascading set of failures where AWS cannot delete the next set of resources because they still have active dependencies on the previous set.

Sample error (from EventsView Logs):

```
Failed to delete cluster. Reason: Failed to wait for deletion: DELETE_FAILED : msg: failed to delete aws stack
Cloudformation says resource xyz has a dependent object (Service: AmazonEC2; Status Code: 400; Error Code: DependencyViolation; Request ID: 815928e2-277e-4b8b-9fed-4b89716a205b) EKS - cluster still existed, was blocking CF delete
```

Cloudera Machine Learning includes a Force Delete option now that will remove the workspace from the Cloudera Machine Learning service. However, this not mean all the underlying resources have been cleaned up. This is where tags are very useful.

If you assigned tags to the workspace at the time of provisioning, you can use the AWS console or the CLI to query the tags associated with the workspace to see if any resources need to be cleaned up manually. Tags associated with a workspace are available on the workspace Details page.

You can search by tags in the EC2 and VPC services. You can also use the AWS CLI to search for specific tags: [resourcegroupstaggingapi](#)

Users unable to access provisioned Cloudera Machine Learning Workspaces

If you have provisioned a workspace but your colleagues cannot automatically access the workspace using Cloudera Data Platform Single-Sign on, make sure that you have completed all the steps required to grant users access to workspaces: [Configuring User Access to Cloudera Machine Learning](#). All Cloudera Machine Learning users must have Cloudera Data Platform accounts.

Troubleshooting Cloudera Machine Learning Workspaces on AWS

This topic describes the Cloudera Machine Learning Workspace provisioning workflow and tells you how to start debugging issues with Cloudera Machine Learning Workspaces on AWS.

Cloudera Machine Learning Workspace Provisioning Workflow

When you provision a new Cloudera Machine Learning Workspace on AWS, Cloudera Machine Learning performs the following actions:

1. Communicates with the Cloudera Management Console to check your AWS credentials. It will also enable Single Sign-On so that authorized Cloudera Data Platform users are automatically logged in to the workspace that will be created.
2. Provisions an NFS filesystem for the workspace on your cloud service provider. On AWS, Cloudera Machine Learning will provision storage on EFS.
3. Provisions a Kubernetes cluster on your cloud service provider. This cluster runs the workspace infrastructure and compute resources. On AWS, Cloudera Machine Learning provisions an EKS cluster.
4. Mounts the provisioned NFS filesystem to the Kubernetes cluster.
5. Provisions TLS certificates for the workspace using LetsEncrypt.
6. Registers the workspace with the cloud provider's DNS service. On AWS, this is Route53.
7. Installs Cloudera Machine Learning onto the EKS cluster.

Troubleshooting resources

Any of the steps listed above can experience failures. To start debugging, you will require access to one or more of the following resources.

- [Workspace > Details Page](#)

Each workspace has an associated Details page that lists important information about the workspace. To access this page, sign in to Cloudera Data Platform, go to Cloudera Machine Learning Workspaces and click on the workspace name.

This page lists basic information about the workspace such as who created it and when. More importantly, it includes a link to the environment where the workspace was created, a link to the underlying EKS cluster on AWS, a list of tags associated with the workspace, and the computing resources in use. The rest of this topic explains how to use these resources.

- [Workspace > Events Page](#)

Each workspace also has an associated Events page that captures every action performed on the workspace. This includes creating, upgrading, and removing the workspace, among other actions. To access this page, sign in to Cloudera Data Platform, go to Cloudera Machine Learning Workspaces, click on the workspace name, and then click Events.

Click the View Logs button associated with an action to see a high-level overview of all the steps performed by Cloudera Machine Learning to complete the action.

The Request ID associated with each action is especially useful in case of a failure as it allows Cloudera Support to efficiently track the series of operations that led to the failure.

- [Environment > Summary Page](#)

Cloudera Machine Learning Workspaces depend quite a bit on the Cloudera Machine Learning Workspace environment in which they are provisioned. Each environment's Summary page lists useful information that can help you debug

issues with the Cloudera Machine Learning service. You can access the environment directly from the workspace Details page.

This page includes important information such as:

- **Credential Setup** - Tells you how security has been configured for the environment. Your cloud credential gives Cloudera Data Platform access to the region and virtual network that make up the environment thus allowing Cloudera Data Platform to provision resources within that environment.
- **Region** - The AWS region where the environment is provisioned. This is especially important because it tells you which region's AWS console you might need to access for further debugging.
- **Network** - The VPC and subnets that were created for the environment. Each Cloudera Machine Learning Workspace requires a set of unique IP addresses to run all of its associated Kubernetes services. If you begin to run out of IP addresses, you will need these VPC and subnet IDs to debug further in the AWS console.
- **Logs** - When you create a Cloudera Data Platform environment, you are asked to specify an S3 bucket in that environment that will be used to store logs. All Cloudera Machine Learning operational logs and Spark logs are also written to this bucket.

You can use the AWS console to access these logs. Alternatively, Site administrators can download these logs directly from their workspace Site Admin panel (Admin Support).



Note: If you file a support case, Cloudera Support will not automatically have access to these logs because they live in your environment.

- **AWS Management Console**

If you have all the relevant information about the environment and the workspace, you can go to the AWS console (for the region where your environment was created) to investigate further. The AWS Management Console has links to dashboards for all the services used by Cloudera Machine Learning.

- **EC2**

You can use the EC2 service dashboards to check the instance-type (CPU, GPU), VPC, subnet, and security group limits imposed on your AWS account. For example, there is typically a limit of 5 VPCs per region.

If you need more resources, submit a request to Amazon to raise the limit of a resource.

- **EKS**

EKS will give you more information such as the version of Kubernetes Cloudera Machine Learning is using, network information, and the status of the cluster. The workspace Details page gives you a direct link to the provisioned EKS cluster on the AWS console.



Note: By default, users do not have Kubernetes-level access to the EKS cluster. If a user wants to use kubectl to debug issues with the EKS cluster directly, an MLAdmin must explicitly grant access using the instructions provided here: [Granting Remote Access to Cloudera Machine Learning Workspaces on EKS](#).

- **VPC**

Use the VPC ID obtained from the Cloudera Data Platform environment Summary page to search for the relevant VPC where you have provisioned or are trying to provision an Cloudera Machine Learning Workspace. Each Cloudera Machine Learning Workspace requires a set of unique IP addresses to run all of its associated Kubernetes services. You can use this service to see how many IP addresses are available for each subnet.

- **S3**

Use the S3 bucket configured for the environment to check/download logs for more debugging.

- **Tags**

When provisioning an Cloudera Machine Learning Workspace, you will have the option to assign one or more tags to the workspace. These tags are then applied to all the underlying AWS resources used by the workspace. If failures occur during provisioning or de-provisioning, it can be very useful to simply query the

tags associated with the workspace to see if any resources need to be cleaned up manually. Tags associated with a workspace are available on the workspace Details page.

You can search by tags in the EC2 and VPC services. You can also use the AWS CLI to search for specific tags: [resourcegroupstaggingapi](#)

- Trusted Advisor (available with AWS Support)

Use the Trusted Advisor dashboard for a high-level view of how you are doing with your AWS account. The dashboard displays security risks, service limits, and possible areas to optimize resource usage. If you have access to AWS Support, Cloudera recommends to review your current account status with Trusted Advisor before you start provisioning Cloudera Machine Learning Workspaces.

Troubleshooting Cloudera Machine Learning Workspaces on Azure

You can collect logs to troubleshoot issues that occur in Cloudera Machine Learning Workspaces with Azure.

How to access Azure logs

Logs from the AKS control plane can be found in the "Logs" blade of the liftie-xxxxxxx resource group (not to be confused with the "Logs" blade of the AKS cluster itself or the Log Analytics Workspace in that resource group). The logs can be looked up using a [query language](#) developed by Microsoft.

Cluster fails to scale down

If a worker node is idle but is not being scaled down, check the cluster autoscaler logs.

Use this example to look up the logs:

```
AzureDiagnostics | where Category ==  
    "cluster-autoscaler"
```

The logs list the pods that are scheduled on a given node that are preventing it from being scaled down, or other reasons for its scaling decisions. Services running in the kube-system namespace (such as tunnelfront, or metrics-server) have been known to delay scale-down when scheduled on an otherwise idle node.

Delete Cloudera Machine Learning Workspace fails

If you delete a workspace, and the delete operation fails, you can use Force delete to remove the workspace.

In this case, Cloudera Machine Learning attempts to delete associated cloud resources for the workspace including metadata files. However, users should check that all such resources have been deleted, and delete manually if necessary.

Logs for Cloudera Machine Learning Workspaces

You can access logs to troubleshoot issues with the Cloudera Machine Learning service and your workloads on Cloudera Machine Learning Workspaces.

Access to logs

When you create a Cloudera Data Platform environment, you specify an S3 bucket (on AWS) or an Azure Storage container (on Azure) in that environment for storing logs. If you have access to the log storage, you can use the AWS or Azure console to access certain Cloudera Machine Learning and Spark logs directly. You can get the details of the specific bucket or container from the [Troubleshooting Cloudera Machine Learning Workspaces on AWS](#) page for the environment.



Note: If you file a support case, Cloudera Support will not automatically have access to these logs because they live in your environment.

Cloudera Machine Learning Workspace access to logs

Cloudera Machine Learning Workspace users also have access to these logs depending on their authorization level:

- Site Administrators

Site administrators can download the same logs directly from their workspace Site Admin panel (Admin Support). For more details, see [Downloading diagnostic bundles for a workspace](#) on page 22.

- Data Scientists

While data scientists do not have access to the full set of workspace logs, they do have access to engine logs for their own workloads (sessions/jobs/experiments). While in an interactive session or on a job/experiment's Overview page, click Download Logs at any time to review the full set of logs for that workload's engine. In the case of Spark workloads, Spark executor and event logs are also downloaded as part of this bundle.

Related Information

[Configure lifecycle management for logs on AWS](#)

[Configure lifecycle management for logs on Azure](#)

Downloading diagnostic bundles for a workspace

This topic describes how to download diagnostic bundles for an Cloudera Machine Learning Workspace.

Before you begin

Required Role: MLAdmin

Make sure you are assigned the MLAdmin role in Cloudera Data Platform. Only users with the MLAdmin role will be logged into Cloudera Machine Learning Workspaces with Site Administrator privileges.

Procedure

1. Log in to the web interface.
2. In Cloudera Machine Learning Workspaces, select a workspace.
3. Select Site Administration Support Generate Log Archive .
4. Select the time period from the dropdown.
5. Ensure Include Engines is selected if engine logs are needed (included by default).
6. Select Send to Cloudera to send the diagnostic logs to Cloudera Support.
7. Select Create to generate the logs.
8. When Status is Complete, select Download to download the diagnostics bundles to your machine.

What to do next

The data in the contained bundles may be incomplete. If it does not contain logs for time period you are looking for, there are a number of possible reasons:

- There is a delay between the time the logs are initially generated by a workload and the time they are visible in cloud storage. This may be approximately 1 minute due to buffering during streaming, but can be significantly longer due to eventual consistency in the cloud storage.
- Another user or process may have deleted data from your bucket; this is beyond the control of Cloudera Machine Learning.
- There may be a misconfiguration or an invalid parameter in your request. Retrieving logs requires a valid cloud storage location to be configured for logging, as well as authentication for Cloudera Machine Learning to be set up properly for it. Requests must pertain to a valid engine in a valid project.

Troubleshooting issues with workloads

This section describes some potential issues data scientists might encounter once the Cloudera Machine Learning Workspace is running workloads.

401 Error caused by incompatible Data Lake version

The following error might occur due to an incompatible Data Lake version:

```
org.apache.ranger.raz.hook.s3.RazS3ClientCredentialsException: Exception in
Raz Server;
Check the raz server logs for more details, HttpStatus: 401
```

To avoid this issue, ensure that:

- Data Lake and Runtime (server) version is 7.2.11 or higher.
- Hadoop Runtime add-on (client) used in the Cloudera Machine Learning session is 7.2.11 or higher.
- Spark Runtime add-on version must be CDE 1.13 or higher.

Engines cannot be scheduled due to lack of CPU or memory

A symptom of this is the following error message in the Workbench: "Unschedulable: No node in the cluster currently has enough CPU or memory to run the engine."

Either shut down some running sessions or jobs or provision more hosts for Cloudera Machine Learning.

Workbench prompt flashes red and does not take input

The Workbench prompt flashing red indicates that the session is not currently ready to take input.

Cloudera Machine Learning does not currently support non-REPL interaction. One workaround is to skip the prompt using appropriate command-line arguments. Otherwise, consider using the terminal to answer interactive prompts.

PySpark jobs fail due to Python version mismatch

```
Exception: Python in worker has different version 2.6 than that in driver 2.
7, PySpark cannot run with different minor versions
```

One solution is to install the matching Python 2.7 version on all the cluster hosts. A better solution is to install the Anaconda parcel on all CDH cluster hosts. Cloudera Machine Learning Python engines will use the version of Python included in the Anaconda parcel which ensures Python versions between driver and workers will always match. Any library paths in workloads sent from drivers to workers will also match because Anaconda is present in the same location across all hosts. Once the parcel has been installed, set the PYSARK_PYTHON environment variable in the Cloudera Machine Learning Admin dashboard.

Troubleshooting Spark issues

Consider some common Spark issues and their recommended solutions.

Distributing SparklyR packages

When using SparklyR, you may see an error from Spark similar to the following:

```
ERROR sparklyr: RScript (4922) terminated
```



```
unexpectedly: namespace 'vctrs' 0.6.3 is being loaded, but >= 0.6.4 is required
```

This means you need to upgrade the vctrs package, and configure Spark to use the R libraries available in the .home directory of the Cloudera Machine Learning session. This works because with Spark on Kubernetes, the Spark executors are running in the same cluster and have access to the same underlying filesystem. Include the following code when creating the Spark session:

```
install.packages("vctrs")
print(packageVersion("vctrs")) # expect 0.6.5
...
config$spark.executorEnv.R_LIBS="/home/cdsdw/.local/lib/R/4.3/library"
config$spark.executorEnv.R_LIBS_USER="/home/cdsdw/.local/lib/R/4.3/library"
config$spark.executorEnv.R_LIBS_SITE="/opt/cmladdons/r/libs"
...
spark_apply(test_fx, packages = TRUE) # true by default
```

Troubleshooting Kerberos issues

This topic describes some common Kerberos issues and their recommended solutions.

HDFS commands fail with Kerberos errors even though Kerberos authentication is successful in the web application

If Kerberos authentication is successful in the web application, and the output of klist in the engine reveals a valid-looking TGT, but commands such as `hdfs dfs -ls /` still fail with a Kerberos error, it is possible that your cluster is missing the [Java Cryptography Extension \(JCE\) Unlimited Strength Jurisdiction Policy File](#). The JCE policy file is required when Red Hat uses AES-256 encryption. This library shall be installed on each cluster host and will live under \$JAVA_HOME. For more information, see [Using AES-256 Encryption](#).

Cannot find renewable Kerberos TGT

Cloudera Machine Learning runs its own Kerberos TGT renewer which produces non-renewable TGT. However, this confuses Hadoop's renewer which looks for renewable TGTs. If the Spark 2 logging level is set to WARN or lower, you may see exceptions such as:

```
16/12/24 16:38:40 WARN security.UserGroupInformation: Exception encountered
while running the renewal command. Aborting renew thread. ExitCodeException
exitCode=1: kinit: Resource temporarily unavailable while renewing credential
s

16/12/24 16:41:23 WARN security.UserGroupInformation: PrivilegedActionExcep
tion as:user@CLOUDERA.LOCAL (auth:KERBEROS) cause:javax.security.sasl.SaslEx
ception: GSS initiate failed [Caused by GSSEException: No valid credentials p
rovided (Mechanism level: Failed to find any Kerberos tgt)]
```

This is not a bug. Spark 2 workloads will not be affected by this. Access to Kerberized resources should also work as expected.