Cloudera on Premises Data Services 1.5.5

Cloudera Data Services on premises Resource Utilization

Date published: 2023-12-16
Date modified: 2026-08-26



https://docs.cloudera.com/

Legal Notice

© Cloudera Inc. 2026. 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

Resource Utilization Dashboard.....4

Managing cluster resources using Quota Management.....4

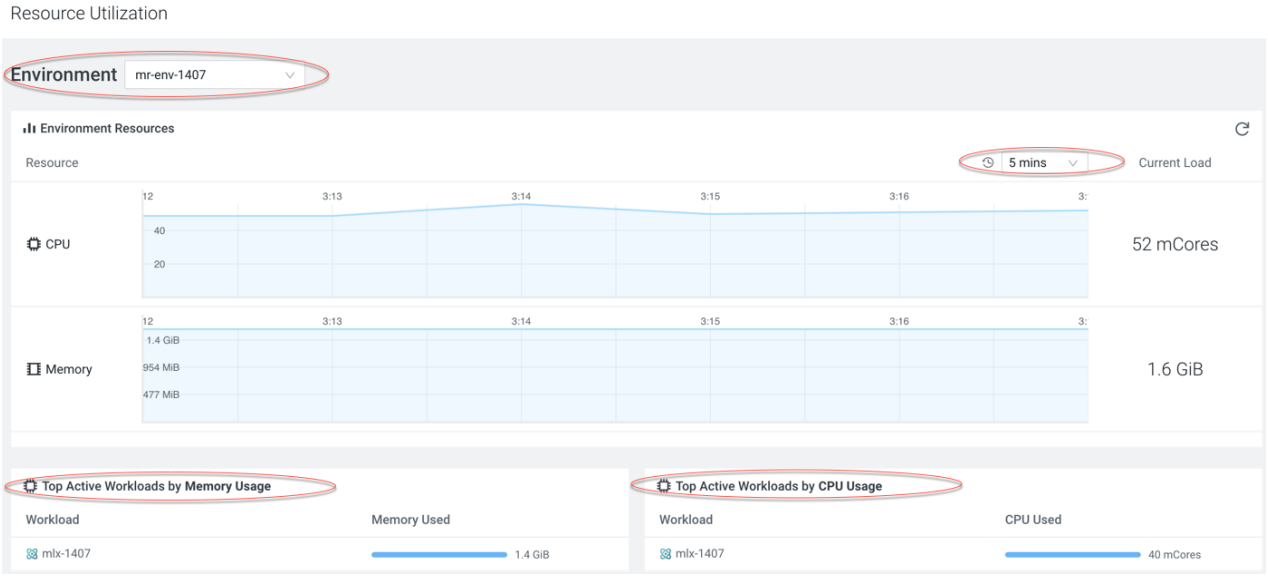
Registering a Cloudera on premises environment..... 10

Resource Utilization Dashboard

Resource Utilization provides an overview of the resources consumed by the Cloudera workloads for a specific environment. You can get information about all compute clusters used across environments and compute resource utilization within these compute clusters. You can view information about the number of cores and the memory used by a specific environment. Using this data, you can plan your resource allocation for the workloads.

You can select a time range from the available options, between five minutes and 15 days, and view resource utilization data of a specific environment.

Note: When you select a time range to display that is longer than the available data, this message appears below each graph. “Note: Some data for the selected time range is missing, which may cause this graph to a ppear misleading. Select a shorter range.”



Managing cluster resources using Quota Management

You can use Quota Management page in the Cloudera Management Console to control how cluster resources are allocated across all Cloudera Data Services on premises.

Using Resource pools in a Cloudera Data Services on premises

Quotas are managed and set up in a hierarchical structure. A single root level is created for the whole cluster and cannot be manipulated. The second level in the structure is the environment node. Multiple environments can be created in a single cluster. For more details, see Registering a Cloudera on premises environment on page 10.

Quotas are primarily set up and managed through the data services. Each data service provides its own integration. When adding a Data Service to an environment, a fixed base pool for the Data Service is added. A data service is always part of an environment and counts towards the quota set for the environment.

The fixed base pool is the standard node for adding services.

- For information about specifying pools in the Cloudera Data Engineering service, see Adding a Cloudera Data Engineering service and Creating virtual clusters .
- For information about specifying pools in the Cloudera AI service, see Configuring Quotas.
- For information about specifying pools in the Cloudera Data Warehouse service, see Resource templates for CDW Private Cloud pods.

Environment creation

While creating an environment a pool is created for that environment. The quota for the environment can be set as part of the flow of creating an environment and adjusted if needed using the Resource Management UI that is part of the console.

The recommendation is to set a quota even if only one environment will be created. All services created within an environment cannot use more resources than the quota set for the environment. Minimum requirements for the services are documented as part of the data service.

The monitoring service that is automatically installed as part of the environment creation requires 30 GB memory and 4 cpu resources. Setting a quota on the environment that is smaller than the monitoring requirements will cause failures in the environment creation.

Upgrade from earlier releases to 1.5.5

As part of the upgrade process the pool structure will be migrated to the new standard structure. Pools that were created by earlier versions will be moved around to align with the new structure.

As part of these changes the following updates will be made:

- The "default" node is no longer used and will be removed.
- A new level for environments will be inserted replacing the existing "default" node.
- A new level for services will be inserted using fixed names for the services: "cde", "cml" and "cdw"
- The option to choose a base pool in the hierarchy as part of the service install will be removed.

Even if no quota management was used in the cluster the environment level will be created for all existing environments. The environment nodes will not have a quota set as part of the migration. The services within each environment are allowed to use the whole cluster. This is equivalent to the pre 1.5.5 behaviour.

Existing service related pools with their quota setting will be retained if the pool before and after upgrade exists. Newly created service related pools will be created without a quota, like the environment level pools.



Note:

When quota management is turned ON and a quota is decreased to below the current usage, no new workloads will be scheduled. The workloads will stay pending, waiting for resources to become available. If the running workloads do not release any resources, new workloads will be blocked. Therefore, when lowering a quota, it is important to consider the current usage. If the quota is set below the current usage, the administrator must manually intervene by evicting workloads to avoid pending workloads becoming permanently blocked.

Node setup and grouping

Quota management enables you to control how resources are allocated within your Cloudera Data Services on premises clusters. In order to prevent a single workload from consuming all available cluster resources, you can limit the number of CPUs, GPUs, and memory allocated by application, user, business units, or Data Service by defining resource pools that define resource limits. Pools are organized in a hierarchical manner by defining nodes in the hierarchy with resource limits, which can then be subdivided as needed to allocate resources for an organization and to allocate resources to cluster or environment wide services such as the monitoring service.

Editing the hierarchy or nodes in the hierarchy is restricted at certain levels in the UI. These levels are directly managed by the services:

- Level 1: root level: pools cannot be edited, added or removed at this level
- Level 2: environment level: quotas and settings can be edited, pools cannot be added or removed at this level
- Level 3: service level: quotas and settings can be edited, pools cannot be added or removed at this level



Note: Levels 4 and below are not restricted.

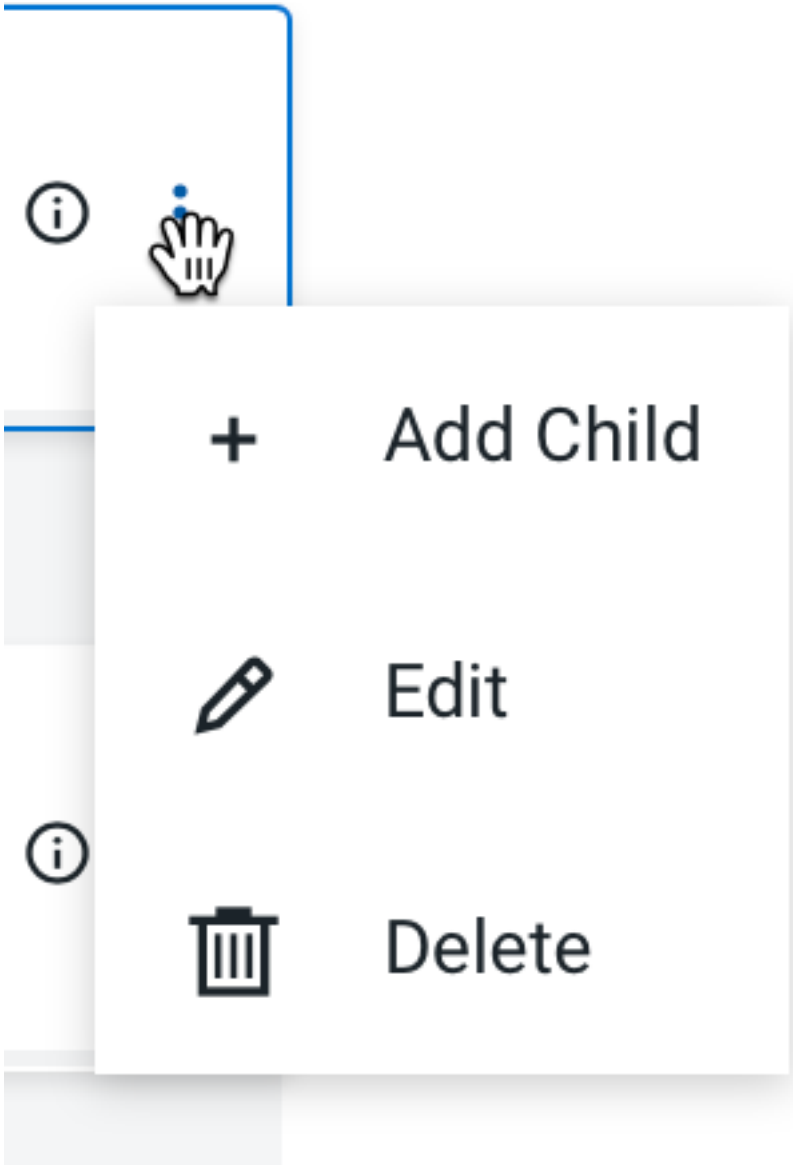
Accessing the Quotas Page

Important: Any quotas created with the Cloudera Data Engineering Service or Virtual Cluster names must not be edited using the Quota Management UI. To update these quotas, use the Cloudera Data Engineering UI.

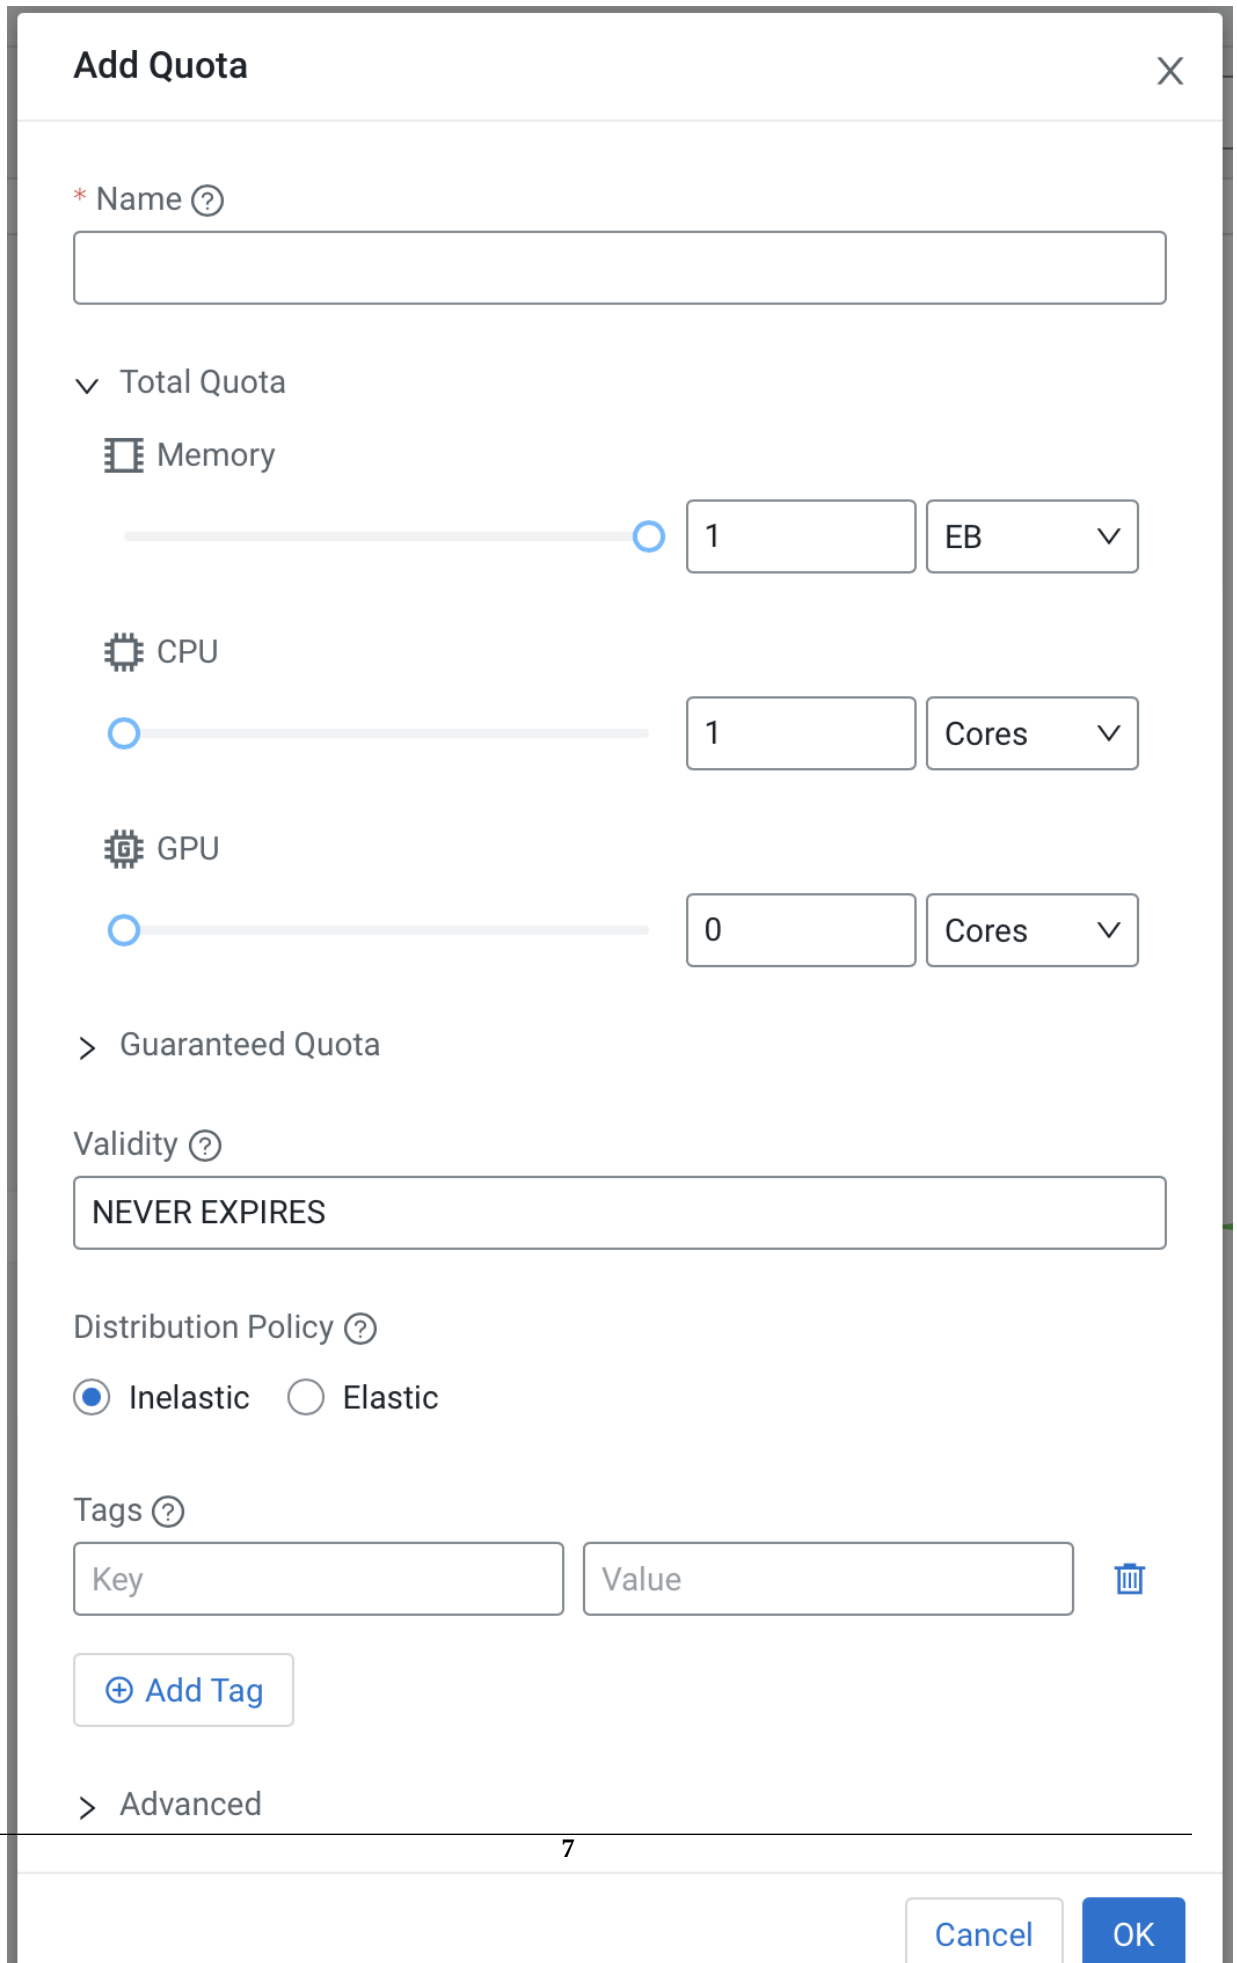
1. Open the Cloudera console.
2. Go to the Cloudera Management Console service.
3. Click Resource Utilization in the left navigation panel.
4. Select the Quotas tab.

Creating child nodes (resource pools),

1. Click the actions menu in a pool and select Add Child.



The Add Quota dialog box displays.


Add Quota
* Name
Total Quota
Memory
1
EB
CPU
1
Cores
GPU
0
Cores
Guaranteed Quota
Validity
NEVER EXPIRES
Distribution Policy
Inelastic
Elastic
Tags
Key
Value
Add Tag
Advanced
Cancel
OK

2. Enter the following information:
   - Name – Enter a name for this resource pool.
   - Memory – Use the sliders to select the memory allocation for the pool. Use the drop-down list to select the units.
   - CPU – Use the sliders to select the CPU allocation for the pool. You can choose how CPUs are counted by selecting Cores or Millicores from the drop-down list.
   - GPU – Use the sliders to select the GPU allocation for the pool.
   - Validity – Enter the length of time that the pool can remain active. To keep the pool active indefinitely, enter -1.
   - Distribution policy – The distribution type of the parent quota to the child pools (elastic or inelastic).
   - Tags – Tags provide a way to add user-defined name/value pairs as metadata for the pools. Tags are not currently used in this release.



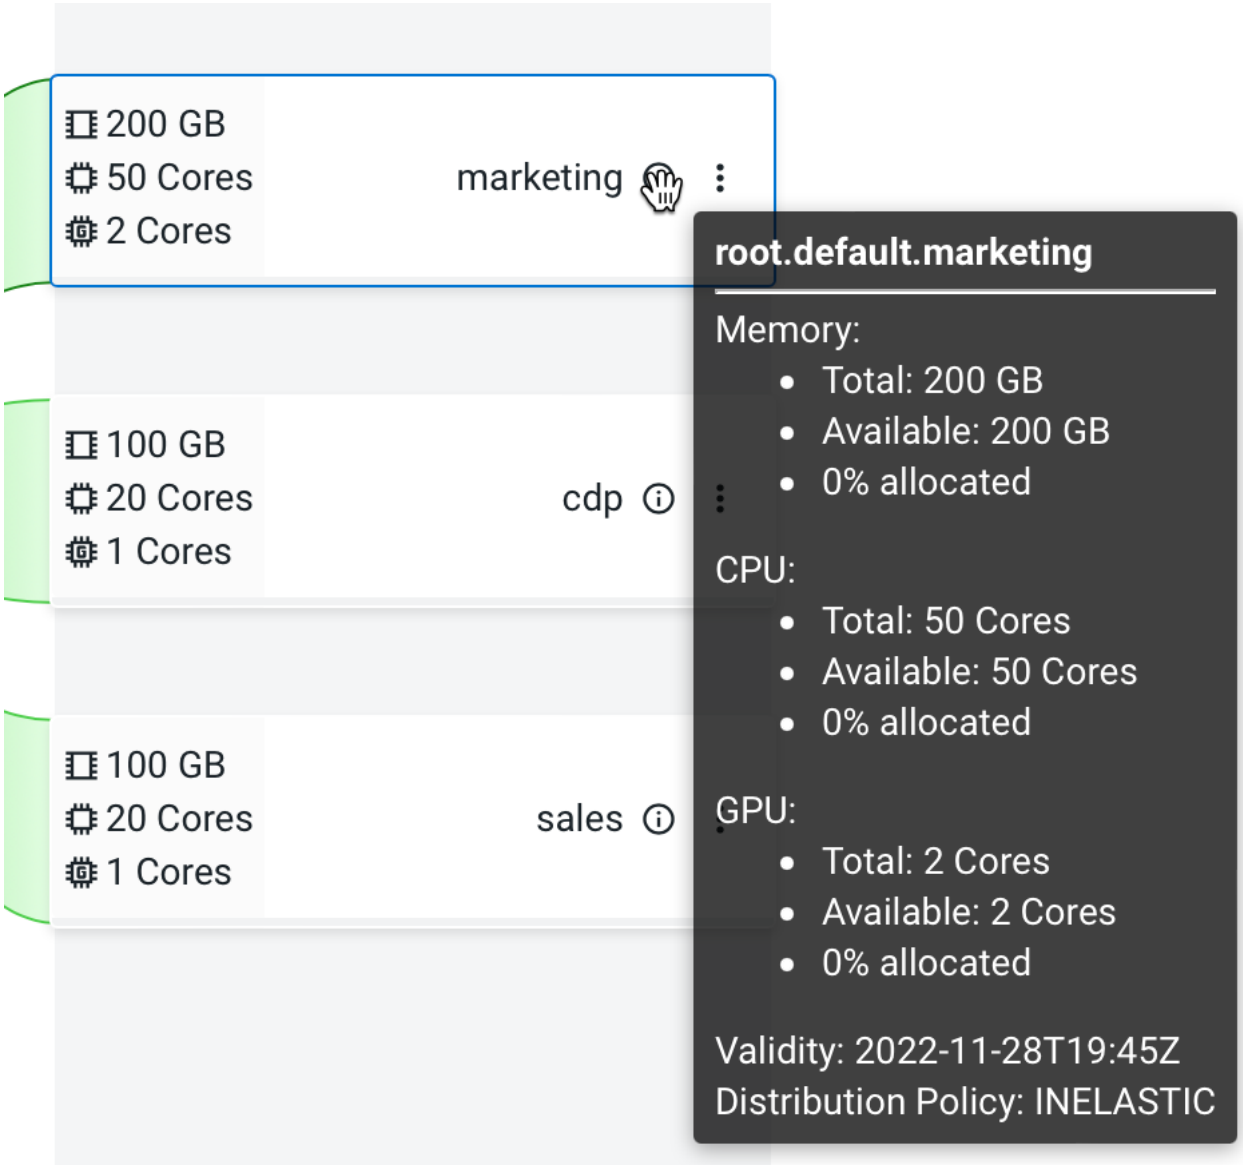
Note: The maximum value of the sliders is limited by the resources of the parent pool and the amount of resources not used by other pools.

Editing a resource pool

You can edit a resource pool after creating it by clicking the actions menu for the pool and selecting Edit.

Viewing details of a resource pool

You can view the resource allocations and other information by clicking the “i” icon. The namespace for the pool appears at the top:



Sorting the resource pool display

Click the Sort by drop down list and select a sorting option to sort the resource pools within each level.

Hierarchical quotas and guaranteed resources

The quota and guaranteed resources are defined in the hierarchy of pools. Rules apply to what can be set as a quota or guaranteed resource for a pool.

The top level pool, the root, does not allow setting a quota or guaranteed resources. The quota and the guaranteed resources will be reflect the cluster size during the quota checks.

For any pool in the hierarchy the following rules apply:

1. Quota must be the same size or larger than the guaranteed resources set on a pool

   Example: a guaranteed resource of memory = 1 TB is not allowed if quota is set to memory = 100 GB
2. Quota set on a pool cannot be larger than the quota set on its parent pool

   Example: a child pool quota of memory = 1 TB is not allowed if parent pool quota is set to memory = 100 GB
3. Guaranteed resource set on a pool cannot be larger than the guaranteed resource of its parent pool

Example: a guaranteed resource of CPU = 10 on the child is not allowed if guaranteed of the parent is set to CPU = 5

4. The sum of all guaranteed resources of pools belonging to the same parent pool cannot be larger than the guaranteed resource of that parent pool.

   Example: two child pools with a guaranteed resource of CPU = 10 each is not allowed if guaranteed of the parent is set to CPU = 15

Note: Levels 4 and below are not directly managed by a service and can be edited, added and deleted. Limitation as described for quotas and guarantees in the hierarchy still apply.

Note: Since the root pool does not have a quota or guaranteed resource set the pools below it, the environment level pools, can be configured to have a quota or guarantee larger than the cluster. This does not mean the pool can grow to a size larger than the what is available in the cluster.

If an inelastic quota distribution policy is used for a pool some extra rules apply:

1. Quotas set on a child pool must specify at least the same set of resource types specified on the parent pool.

   Example: parent specifies memory and CPU as types, the child must specify at least memory and CPU in its quota. It may add others.
2. The sum of all quotas of pools belonging to the same parent pool cannot be larger than the quota of that parent pool.

   Example: two child pools with a quota of CPU = 10 each is not allowed if quota of the parent is set to CPU = 15

Registering a Cloudera on premises environment

After you have set up the Cloudera on premises requirements, you can register the environment.

Steps

Management Console UI

1. Sign into the Cloudera console.
2. Click Environments.
3. On the Environments page, click Register Environment.
4. On the Register Environment page, provide the required information.

   Environment

   Environments in Cloudera on premises provide shared data, security, and governance (metadata) for your Cloudera AI and Cloudera Data Warehouse applications.

Property	Description
Environment Name	Enter a name for your environment. This name will be used to refer to this environment in Cloudera. Note: Cloudera Data Warehouse service requires that you specify the environment name less than 45 characters long. This is because Cloudera Data Warehouse uses a deterministic namespace and adds a prefix to the environment name. The length of the namespace ID after Cloudera Data Warehouse applies a prefix to the Environment name, including the hyphen (-), should not exceed 63 characters.

   Resource Quota

   Optionally you can set quota for CPU, Memory, and GPU resources for this environment. When setting a quota for the environment, ensure that it satisfies the resource requirement of the data services to be created within the environment. However, if you do not specify a quota for a resource, then no quota will be set for that resource at the time of the environment creation.

Note: Ensure that at least 4 CPU Cores and 30 GB of RAM are reserved for Monitoring.

Property	Description
CPU (Cores)	You can specify a CPU quota (in Cores) for the environment. The quota must be a positive number.
Memory (GB)	You can specify a Memory quota (in GigaBytes) for the environment. The quota must be a positive number.
GPU (Cores)	You can specify a GPU quota (in Cores) for the environment. The quota must be a positive number.

Compute Cluster Resources

To run workloads, you must specify a Kubeconfig file to register a Kubernetes cluster with Cloudera on premises.

Property	Description
Kubernetes Configurations	Click Upload Files, then select a Kubeconfig file to enable Cloudera to access a Kubernetes cluster.
Storage class	The storage class on the OpenShift cluster. If you do not specify this value, the default storage class is used.
Domain	The default domain suffix for workload applications.

Data Lake

A Data Lake refers to the shared security and governance services in a Cloudera Data Center cluster linked to a Cloudera on premises environment, and managed by Cloudera Manager. To register an environment, Cloudera on premises needs to access Cloudera Manager and its Data Lake services.

Note: When registering an environment in Cloudera Data Services on premises - Data Lake services - Cloudera Manager, you need to use FQDN (Fully Qualified Domain Name) rather than IP address.

Parameter	Description
Cloudera Manager URL	The Cloudera Manager URL.
Cloudera Manager Admin Username	The Cloudera Manager administrator user name.
Cloudera Manager Admin Password	The Cloudera Manager administrator password.

5. Under Data Lake, click Connect.

   When Cloudera on premises has successfully connected to Cloudera Manager, a confirmation message appears, along with the Data Lake cluster services.
6. Click Register.

   The environment page appears. The new environment is also listed on the Environments page.

Note: You can repeat steps 2 through 5 to register more environments. If required, you can select a different Data Lake cluster while registering the additional environment. Cloudera Data Services on premises has multiple base cluster support, and users can create multiple environments each pointing to a different base or data lake cluster.

CLI

You can use the following command to create a new environment:

```
cdp environments create-private-environment \
--environment-name <value> \
--address <value> \
--authentication-token <value> \
----cluster-names <value>
```

For a detailed description of the command properties, use cdp environments create-private-environment --help

Creating multiple environments with different base or Data Lake clusters

To register an environment with a data lake cluster managed by a Cloudera Manager that is different from your existing Cloudera Manager, you need to add the certificates of the new Cloudera Manager to the Cloudera Management Console. If the existing Cloudera Manager and the new Cloudera Manager share the same root CA, and the root CA is already uploaded as the data lake certificate, then no additional certificate needs to be added.

Cloudera Manager certificates can be accessed from the Cloudera Manager server.

1. To get the certificate path on the Cloudera Manager server, navigate to the Administration tab on your Cloudera Manager UI.
2. Select Settings, and within the Settings, select Security as the category.
3. In the list of settings, Cloudera Manager TLS/SSL Server Keystore File Location points to the Cloudera Manager server certificate. This certificate needs to be combined with the certificates of your existing Cloudera Manager server(s).

   Note: The combined certificate needs to be in X.509 PEM format.

   For information on combining certificates, see Configuring multiple base clusters with ECS .
4. The combined certificate then needs to be uploaded as a Datalake certificate to Cloudera Management Console > Administration > CA Certificates . This ensures that the Cloudera Data Services on premises cluster is configured to support both the new and the existing data lake clusters.