

CDP Private Cloud Experiences Resource Utilization

Date published: 2020-12-16

Date modified: 2023-06-08

CLUSTERA

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

Resource Utilization Dashboard.....	4
Managing cluster resources using Quota Management.....	4

Resource Utilization Dashboard


Resource Utilization provides an overview of the resources consumed by the CDP 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 appear misleading. Select a shorter range.”

Resource Utilization



Managing cluster resources using Quota Management

You can use Quota Management page in the Management Console to control how cluster resources are allocated across all CDP Private Cloud Data Services.

Overview


Quota management enables you to control how resources are allocated within your CDP Private Cloud Data Services 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-wide services such as the Control Plane. You can also use this feature to audit workload ownership and resource usage.

A system administrator defines the initial groupings, starting with the default tenant node and first creating child nodes that define quotas for CDP system resources. You will need to edit the default node settings to reflect the current total resources of your cluster.


Add Quota ✕

* Name ?


✓ Total Quota

 Memory

1 EB ▼

 CPU

1 Cores ▼

 GPU

0 Cores ▼


> Guaranteed Quota

Validity ?

Distribution Policy ?

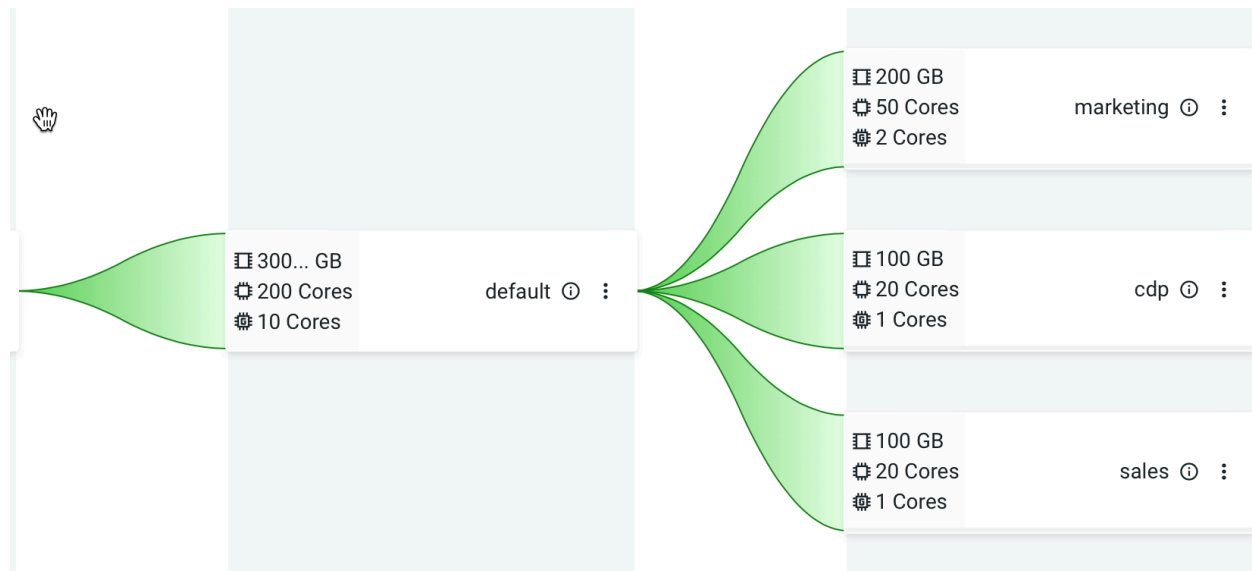
Inelastic Elastic

Tags ?

Key	Value	

> Advanced 5

Next the administrator creates one or more child nodes, in this example, nodes are created for the marketing and sales business units and a node is created for CDP resources.



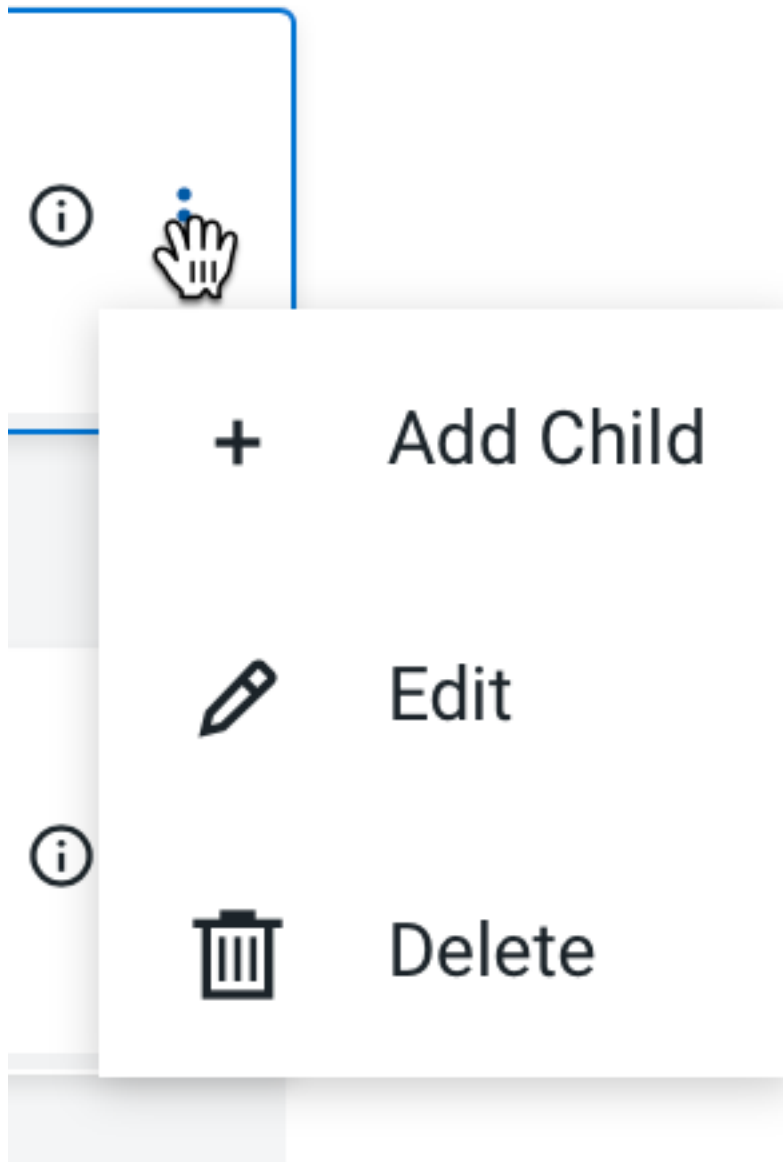
When Data Service users create new artifacts in their service's interface they can reference a parent pool name in which to run their workload. The resources available for the workload are limited by the quotas established and consumed in the parent for pool. For example, the Parent pool, root.default.marketing refers to the "marketing" business unit pool node defined in the hierarchy.

Accessing the Quotas Page

1. Open the CDP console.
2. Go to the 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

▼

 CPU

▼

 GPU

▼


> Guaranteed Quota

Validity ?

Distribution Policy ?

Inelastic Elastic

Tags ?

Key	Value	
<input type="text"/>	<input type="text"/>	

> Advanced

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 that 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:

The screenshot displays a list of three resource pools. The first pool, named 'marketing', has 200 GB of memory, 50 Cores of CPU, and 2 Cores of GPU. The second pool, named 'cdp', has 100 GB of memory, 20 Cores of CPU, and 1 Core of GPU. The third pool, named 'sales', has 100 GB of memory, 20 Cores of CPU, and 1 Core of GPU. A detailed view for the 'marketing' pool is shown on the right, titled 'root.default.marketing'. This view provides the following information:

- Memory:**
 - Total: 200 GB
 - Available: 200 GB
 - 0% allocated
- CPU:**
 - Total: 50 Cores
 - Available: 50 Cores
 - 0% allocated
- GPU:**
 - Total: 2 Cores
 - Available: 2 Cores
 - 0% allocated

Additional details for the 'marketing' pool include a validity period of 2022-11-28T19:45Z and a distribution policy of INELASTIC.

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.

Using Resource pools in a Data Service

When creating artifacts in a Data Service, you specify which pool to use to allocate resources to the artifacts.

- For information about specifying pools in the Cloudera Data Engineering (CDE) service, see [Adding a Cloudera Data Engineering service](#) and [Creating virtual clusters](#).
- For information about specifying pools in the Cloudera Machine Learning (CML) service, see [Configuring Quotas](#).
- For information about specifying pools in the Cloudera Data Warehouse (CDW) service, see [Enable quota management \(optional\)](#).