

**ModelArts**

# **Resource Management**

**Issue**            01  
**Date**             2024-04-30



**Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2024. All rights reserved.**

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Cloud Computing Technologies Co., Ltd.

## **Trademarks and Permissions**



HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

## **Notice**

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

---

# Contents

---

|  |           |
|--|-----------|
| <b>1 Resource Pool</b>   | <b>1</b>  |
| <b>2 Elastic Cluster</b>   | <b>3</b>  |
| 2.1 Comprehensive Upgrades to ModelArts Resource Pool Management Functions | 3         |
| 2.2 Creating a Resource Pool   | 5         |
| 2.3 Viewing Details About a Resource Pool                                  | 10        |
| 2.4 Resizing a Resource Pool   | 15        |
| 2.5 Setting a Renewal Policy   | 17        |
| 2.6 Modifying the Expiration Policy  | 18        |
| 2.7 Migrating the Workspace  | 18        |
| 2.8 Changing Job Types Supported by a Resource Pool                        | 20        |
| 2.9 Upgrading a Resource Pool Driver                                       | 21        |
| 2.10 Deleting a Resource Pool  | 22        |
| 2.11 Abnormal Status of a Dedicated Resource Pool                          | 23        |
| 2.12 ModelArts Network   | 28        |
| 2.13 ModelArts Nodes   | 30        |
| <b>3 Audit Logs</b>  | <b>32</b> |
| 3.1 Key Operations Recorded by CTS   | 32        |
| 3.2 Viewing Audit Logs   | 38        |
| <b>4 Monitoring Resources</b>  | <b>40</b> |
| 4.1 Overview   | 40        |
| 4.2 Using Grafana to View AOM Monitoring Metrics                           | 40        |
| 4.2.1 Procedure  | 40        |
| 4.2.2 Installing and Configuring Grafana                                   | 40        |
| 4.2.2.1 Installing and Configuring Grafana on Windows                      | 41        |
| 4.2.2.2 Installing and Configuring Grafana on Linux                        | 42        |
| 4.2.2.3 Installing and Configuring Grafana on a Notebook Instance          | 44        |
| 4.2.3 Configuring a Grafana Data Source                                    | 48        |
| 4.2.4 Using Grafana to Configure Dashboards and View Metric Data           | 53        |
| 4.3 Viewing All ModelArts Monitoring Metrics on the AOM Console            | 59        |

# 1 Resource Pool

---

## ModelArts Resource Pools

When using ModelArts for AI development, you can use either of the following resource pools:

- **Dedicated resource pool:** It delivers more controllable resources and cannot be shared with other users. Create a dedicated resource pool and select it during AI development. The dedicated resource pool can be an elastic cluster or an elastic BMS.
  - Elastic cluster: It can be Standard or Lite.
    - In a Standard elastic cluster, exclusive computing resources are provided, with which you can deliver instances during job training, model deployment, and environment development on ModelArts.
    - A Lite elastic cluster provides hosted Kubernetes clusters with mainstream AI development plug-ins and acceleration plug-ins for Kubernetes resource users. You can operate the nodes and Kubernetes clusters in the resource pool with provided AI Native resources and tasks.
  - Elastic BMS: It provides different models of xPU BMSs. You can access an elastic BMS through an EIP and install GPU- and NPU-related drivers and software on a specified OS image. To meet the routine training requirements of algorithm engineers, SFS and OBS can be used to store and read data.
- **Public Resource Pool:** provides large-scale public computing clusters, which are allocated based on job parameter settings. Resources are isolated by job. You can use ModelArts public resource pools to deliver training jobs, deploy models, or run DevEnviron instances and will be billed on a pay-per-use basis.

## Differences Between Dedicated Resource Pools and Public Resource Pools

- Dedicated resource pools provide dedicated computing clusters and network resources for users. The dedicated resource pools of different users are physically isolated, while public resource pools are only logically isolated. Compared with public resource pools, dedicated resource pools feature better performance in isolation and security.

- When a dedicated resource pool is used for creating jobs and the resources are sufficient, the jobs will not be queued. When a public resource pool is used for creating jobs, there is a high probability that the jobs will be queued.
- A dedicated resource pool is accessible to your network. All running jobs in the pool can access storage and resources in your network. For example, if you select a dedicated resource pool with an accessible network when creating a training job, you can access SFS data after the training job is created.
- Dedicated resource pools allow you to customize the runtime environment of physical nodes, for example, you can upgrade GPU or Ascend drivers. This function is not supported by public resource pools.

# 2 Elastic Cluster

---

## 2.1 Comprehensive Upgrades to ModelArts Resource Pool Management Functions

ModelArts dedicated resource pools have been upgraded and taken effect at 00:00 GMT+08:00 on March 1, 2023. In the new system, there are only unified ModelArts dedicated resource pools, which are no longer classified as the pools dedicated for development/training and the pools dedicated for service deployment. The new-version dedicated resource pools support flexible configuration of job types, and allow you to manage networks and interconnect VPCs with the networks.

The new dedicated resource pool management page provides more comprehensive functions and displays more information about the resource pools. More details about how to use and manage dedicated resource pools are provided in subsequent sections of this document. If you are new to ModelArts dedicated resource pools, try new-version dedicated resource pools. If you have used ModelArts dedicated resource pools, the old-version pools will be smoothly switched to new-version pools.

Read the following contents to learn about new-version dedicated resource pools.

### Features of New-Version Dedicated Resource Pools

The new-version dedicated resource pool management is a comprehensive technology and product improvement. The main improvements are as follows:

- **Single dedicated resource pool type for diverse purposes:** Dedicated resource pools are no longer classified into those for development/training and those for service deployment. You can run both training and inference workloads in a dedicated resource pool. You can also set the job types supported by a dedicated resource pool based on your needs.
- **Dedicated resource pool network connection:** You can create and manage dedicated resource pool networks on the ModelArts management console. If you need to access resources in your VPC for jobs running in a dedicated resource pool, interconnect the VPC with the dedicated resource pool network.

- **More cluster details:** The new-version dedicated resource pool details page provides more cluster details, such as jobs, nodes, and resource monitoring, helping you learn about the cluster status and better plan and use resources.
- **Cluster GPU/NPU driver management:** On the new-version dedicated resource pool details page, you can select an accelerator card driver and perform change upon submission or smooth upgrade of the driver based on service requirements.
- **Fine-grained resource allocation (coming soon):** You can divide your dedicated resource pool into multiple small pools and assign different quotas and permissions to each small pool for flexible and refined resource allocation and management.

More features will be provided in later versions for a better user experience.

## Can I Continue to Use the Existing Dedicated Resource Pools After the Upgrade Takes Effect?

If you have created dedicated resource pools, you can still access the old-version dedicated resource pool (elastic cluster) management page on the ModelArts management console and use the created resource pools, but you cannot create dedicated resource pool on that page. ModelArts allows you to migrate existing dedicated resource pools to the new management page. You will be contacted to complete the migration and this does not require you to perform any operations. In addition, the migration does not affect the workloads running in the dedicated resource pools. Pay attention to the easy-to-use new management functions of dedicated resource pools. There is no change in creating training jobs or inference services.

## Will New-Version Dedicated Resource Pools Be More Expensive?

The charging unit and unit price of the new-version dedicated resource pools are the same as those of the old-version dedicated resource pools. If you do not scale in or out your dedicated resource pools, the fee will not change. In addition, more value-added features, such as subpool division, elastic sharing, and data acceleration, will be provided in later versions to better allocate compute resources and improve cost-effectiveness.

## Differences Between New and Old Dedicated Resource Pools

- In the old version, the dedicated resource pools dedicated for development/training are separated from those dedicated for service deployment. In addition, the pools of the two types offer different functions and their user experience varies. In the new version, the dedicated resource pools of the two types are unified. You only need to configure one or multiple job types. Then, the dedicated resource pool automatically supports the configured job type.
- New dedicated resource pools inherit all functions of the old ones and have greatly improved user experience in key functions such as purchasing and resizing a resource pool. Use new dedicated resource pools for smooth, transparent experience.
- Additionally, the new dedicated resource pools offer enhanced functions, for example, allowing you to upgrade GPU or Ascend drivers, view details about job queuing, and use one network for multiple pools. More new functions of the new dedicated resource pools are coming soon.

## How Can I Get Help or Provide Feedback if I Encounter Problems During Use?

Similar to other ModelArts functions, you can report problems or obtain help in the sidebar of the console. In addition, you are advised to read the subsequent sections of this document to further understand how to use ModelArts dedicated resource pools. Submit a service ticket for more requirements.

### Instructions of Dedicated Resource Pools

- If you use dedicated resource pools for the first time, get started by reading [Resource Pool](#).
- Create a dedicated resource pool by referring to [Creating a Resource Pool](#).
- View the details about a created dedicated resource pool by referring to [Viewing Details About a Resource Pool](#).
- If the specifications of a dedicated resource pool do not meet your service requirements, adjust the specifications by referring to [Resizing a Resource Pool](#).
- Set or change job types supported by a dedicated resource pool by referring to [Changing Job Types Supported by a Resource Pool](#).
- Upgrade the GPU/Ascend driver of your dedicated resource pools by referring to [Upgrading a Resource Pool Driver](#).
- If a dedicated resource pool is no longer needed, delete it by referring to [Deleting a Resource Pool](#).
- If any exception occurs when you use a dedicated resource pool, handle the exception by referring to [Abnormal Status of a Dedicated Resource Pool](#).
- Manage dedicated resource pool networks or interconnect VPCs with the networks by referring to [ModelArts Network](#).

## 2.2 Creating a Resource Pool

This section describes how to create a dedicated resource pool.

### Procedure

1. Log in to the ModelArts console. In the navigation pane, choose **Dedicated Resource Pools > Elastic Cluster**.

#### NOTE

For new users, only new-version elastic clusters are available on the ModelArts console. For users who have used old-version dedicated resource pools, they can access both old-version and new-version elastic clusters.

2. On the **Resource Pools** tab, click **Create** and configure parameters.



**Table 2-1** Dedicated resource pool parameters

| Parameter                | Sub-Parameter  | Description   |
|--------------------------|----------------|---|
| Name                     | N/A            | Name of a dedicated resource pool.<br>Only lowercase letters, digits, and hyphens (-) are allowed. The value must start with a lowercase letter and cannot end with a hyphen (-).   |
| Description              | N/A            | Brief description of a dedicated resource pool.   |
| Billing Mode             | N/A            | You can select <b>Pay-per-use</b> .   |
| Resource Pool Type       | N/A            | You can select <b>Physical</b> or <b>Logical</b> . If there is no logical specification, <b>Logical</b> is not displayed.   |
| Job Type                 | N/A            | Select job types supported by the resource pool based on service requirements. <ul style="list-style-type: none"> <li>• <b>Physical: DevEnviron, Training Job, and Inference Service</b> are supported.</li> <li>• <b>Logical: Only Training Job</b> is supported.</li> </ul>   |
| Network                  | N/A            | Network in which the target service instance is deployed. The instance can exchange data with other cloud service resources in the same network.<br>Select a network from the drop-down list box. If no network is available, click <b>Create</b> on the right to create a network. For details about how to create a network, see <a href="#">Creating a Network</a> . |
| Specification Management | Specifications | Select required specifications. Due to system loss, the actual available resources are less than those specified in the specifications. After a dedicated resource pool is created, you can view the actual available resources on the <b>Nodes</b> tab page of the dedicated resource pool details page.   |


| Parameter         | Sub-Parameter          | Description  |
|-------------------|------------------------|--|
|                   | AZ                     | <p>You can select <b>Automatically allocated</b> or <b>Specifies AZ</b>. An AZ is a physical region where resources use independent power supplies and networks. AZs are physically isolated but interconnected over an intranet.</p> <ul style="list-style-type: none"> <li>• <b>Automatically allocated</b>: AZs are automatically allocated.</li> <li>• <b>Specifies AZ</b>: Specify AZs for resource pool nodes. To ensure system disaster recovery, deploy all nodes in the same AZ. You can set the number of nodes in an AZ.</li> </ul> |
|                   | Nodes                  | <p>Select the number of nodes in a dedicated resource pool. More nodes mean higher computing performance.</p> <p>If <b>AZ</b> is set to <b>Specifies AZ</b>, you do not need to configure <b>Nodes</b>.</p> <p><b>NOTE</b><br/>It is a good practice to create no more than 30 nodes at a time. Otherwise, the creation may fail due to traffic limiting.</p>  |
|                   | Advanced Configuration | <p>This allows you to set the container engine space.</p> <p>You must enter an integer for the container engine space. It cannot be less than 50 GB, which is the default and minimum value. The maximum value depends on the specifications. To see the valid values, check the console prompt. Customizing the container engine space does not increase costs.</p>   |
| Custom Driver     | N/A                    | This parameter is available only when a GPU or Ascend flavor is selected. Enable this function and select a driver.  |
| GPU Driver        | N/A                    | This parameter is available only when custom driver is enabled. Select a GPU accelerator driver.   |
| Required Duration | N/A                    | Select the time length for which you want to use the resource pool. This parameter is mandatory only when the <b>Yearly/Monthly</b> billing mode is selected.  |
| Auto-renewal      | N/A                    | <p>Specifies whether to enable auto-renewal. This parameter is mandatory only when the <b>Yearly/Monthly</b> billing mode is selected.</p> <ul style="list-style-type: none"> <li>• Monthly subscriptions renew each month.</li> <li>• Yearly subscriptions renew each year.</li> </ul>  |

| Parameter           | Sub-Parameter | Description  |
|---------------------|---------------|--|
| Advanced Options    | N/A           | Select <b>Configure Now</b> to set the tag information, CIDR block, and controller node distribution.  |
| Tags                | N/A           | <p>ModelArts can work with Tag Management Service (TMS). When creating resource-consuming tasks in ModelArts, for example, training jobs, configure tags for these tasks so that ModelArts can use tags to manage resources by group.</p> <p>For details about how to use tags, see <a href="#">How Does ModelArts Use Tags to Manage Resources by Group?</a></p> <p><b>NOTE</b><br/>You can select a predefined TMS tag from the tag drop-down list or customize a tag. Predefined tags are available to all service resources that support tags. Customized tags are available only to the service resources of the user who has created the tags.</p>   |
| CIDR block          | N/A           | <p>You can select <b>Default</b> or <b>Custom</b>.</p> <ul style="list-style-type: none"> <li>• <b>Default:</b> The system randomly allocates an available CIDR block to you, which cannot be modified after the resource pool is created. For commercial use, customize your CIDR block.</li> <li>• <b>Custom:</b> You need to customize K8S container and K8S service CIDR blocks. <ul style="list-style-type: none"> <li>– <b>K8S Container Network:</b> used by the container in a cluster, which determines how many containers there can be in a cluster. The value cannot be changed after the resource pool is created.</li> <li>– <b>K8S Service Network:</b> used when the containers in the same cluster access each other, which determines how many Services there can be. The value cannot be changed after the resource pool is created.</li> </ul> </li> </ul> |
| Master Distribution | N/A           | <p>Distribution locations of controller nodes. You can select <b>Random</b> or <b>Custom</b>.</p> <ul style="list-style-type: none"> <li>• <b>Random:</b> Use the AZs randomly allocated by the system.</li> <li>• <b>Custom:</b> Select AZs for controller nodes.</li> </ul> <p>Distribute controller nodes in different AZs for disaster recovery.</p>   |

3. Click **Next** and confirm the settings. Then, click **Submit** to create the dedicated resource pool.


- After a resource pool is created, its status changes to **Running**. Only when the number of available nodes is greater than 0, tasks can be delivered to this resource pool.

**Figure 2-1** Viewing a resource pool

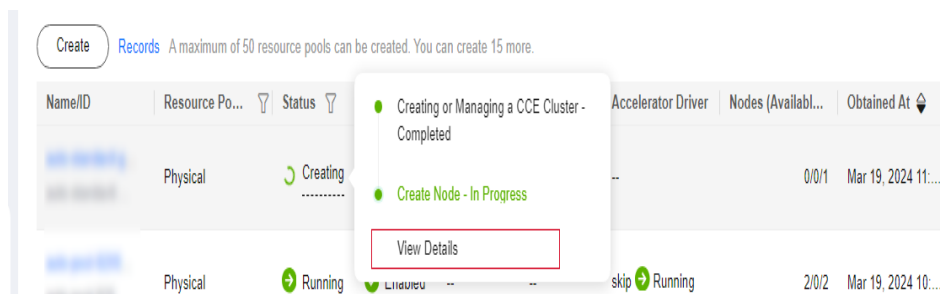


- Hover the cursor over **Creating** to view the details about the creation process. Click **View Details**. The operation record page is displayed.

**Figure 2-2** Creating




**Figure 2-3** Viewing details



- You can view the task records of the resource pool by clicking **Records** in the upper left corner of the resource pool list.

**Figure 2-4** Operation records



**Figure 2-5** Viewing the resource pool status

Records

You can view your order records (excluding logical sub-pools) below. Each record can be retained for a maximum of 90 days.

Enter a name.

| Name/ID                | Operation Status                          | Operation Status | Billing Mode                    | Obtained At                     |
|------------------------|---|------------------|---------------------------------|---------------------------------|
| ^                      | Processing                                | Create           | Pay-per-use                     | Mar 19, 2024 11:38:33 GMT+08:00 |
| Order ID               | --  | Started          | Mar 19, 2024 11:38:33 GMT+08:00 |                                 |
| Initial Specifications | --  | Ended            | --                              |                                 |
| New Specifications     | 1 * modelarts bm gpu 8p100                |                  | Actual Specifications           | --                              |
| Create records         | Project                                   | Status           | Started                         | Ended                           |
|                        | It takes 1 to 10 minutes to manage nod... | Completed        | Mar 19, 2024 11:38:34 GMT+08:00 | Mar 19, 2024 11:45:45 GMT+08:00 |
|                        | Create a node, which takes 10 to 20 mi... | Ongoing          | Mar 19, 2024 11:45:45 GMT+08:00 | --                              |

## FAQs

### What if I choose a flavor for a dedicated resource pool, but get an error message saying no resource is available?

The flavors of dedicated resources change based on real-time availability. Sometimes, you might choose a flavor on the purchase page, but it is sold out before you pay and create the resource pool. This causes the resource pool creation to fail.

You can try a different flavor on the creation page and create the resource pool again.



### Q: Why cannot I use all the CPU resources on a node in a resource pool?

Resource pool nodes have systems and plug-ins installed on them. These take up some CPU resources. For example, if a node has 8 vCPUs, but some of them are used by system components, the available resources will be fewer than 8 vCPUs.

You can check the available CPU resources by clicking the **Nodes** tab on the resource pool details page, before you start a task.

## 2.3 Viewing Details About a Resource Pool

### Resource Pool Details Page

- Log in to the ModelArts console. In the navigation pane, choose **Dedicated Resource Pools > Elastic Cluster**.
- Click  next to the resource pool type or status in the table header. In the top right corner of the list, select **Name** or **Resource ID** to filter resource pools. To obtain the resource ID, go to the **Billing Center > Orders > My Orders** page and click **Details** in the **Operation** column of the target order.
- In the resource pool list, click a resource pool to go to its details page and view its information.
  - If there are multiple resource pools, click  in the top left corner of the details page of one resource pool to switch between resource pools. Click **More** in the top right corner to perform operations such as resize or

delete the resource pool. The available operations vary depending on the resource pool.

- In the **Network** area of **Basic Information**, you can click the number of resource pools associated to view associated resource pools.
- In the extended information area, you can view the monitoring information, jobs, nodes, specifications, and events. For details, see the following section.


## Viewing Jobs in a Resource Pool

On the resource pool details page, click **Jobs**. You can view all jobs running in the resource pool. If a job is queuing, you can view its queuing position.

### NOTE


Only training jobs can be viewed.

Figure 2-6 Jobs



## Viewing Resource Pool Events

On the resource pool details page, click **Events**. You can view all events of the resource pool. The cause of an event is **PoolStatusChange** or **PoolResourcesStatusChange**.

In the event list, click  on the right of **Event Type** to filter events.

- When a resource pool starts to be created or becomes abnormal, the resource pool status changes and the change will be recorded as an event.
- When the number of nodes that are available or abnormal or in the process of being created or deleted changes, the resource pool node status changes and the change will be recorded as an event.


Figure 2-7 Events

| Event Type | Cause                     | Details   | Occurred At                     |
|------------|---------------------------|---|---------------------------------|
| Abnormal   | PoolStatusChange          | Pool status changed, from Running to Abnormal.  | Jan 04, 2024 09:57:32 GMT+08:00 |
| Abnormal   | PoolResourcesStatusChange | Pool resources status changed, available/abnormal/creating/queuing count from 1/1/0/0 to 0/0/0. InetSocketAddress | Jan 04, 2024 09:59:51 GMT+08:00 |
| Normal     | PoolResourcesStatusChange | Pool resources status changed, available/abnormal/creating/queuing count from 1/1/0/0 to 2/0/0. InetSocketAddress | Dec 02, 2023 11:02:44 GMT+08:00 |
| Abnormal   | PoolResourcesStatusChange | Pool resources status changed, available/abnormal/creating/queuing count from 2/0/0 to 1/1/0. InetSocketAddress   | Dec 02, 2023 10:59:03 GMT+08:00 |
| Normal     | PoolResourcesStatusChange | Pool resources status changed, available/abnormal/creating/queuing count from 1/1/0/0 to 2/0/0. InetSocketAddress | Dec 02, 2023 10:14:18 GMT+08:00 |
| Normal     | PoolResourcesStatusChange | Pool resources status changed, available/abnormal/creating/queuing count from 0/0/0 to 1/1/0. InetSocketAddress   | Dec 02, 2023 10:14:16 GMT+08:00 |
| Abnormal   | PoolResourcesStatusChange | Pool resources status changed, available/abnormal/creating/queuing count from 1/1/0/0 to 0/0/0. InetSocketAddress | Dec 02, 2023 10:30:04 GMT+08:00 |
| Abnormal   | PoolResourcesStatusChange | Pool resources status changed, available/abnormal/creating/queuing count from 2/0/0 to 1/1/0. InetSocketAddress   | Dec 02, 2023 10:24:00 GMT+08:00 |

## Viewing Resource Pool Nodes

On the resource pool details page, click **Nodes**. You can view all nodes in the resource pool and the resource usage of each node.

Some resources are reserved for cluster components. Therefore, **CPUs (Available/Total)** does not indicate the number of physical resources on the node. It only displays the number of resources that can be used by services. CPU cores are metered in milicores, and 1000 milicores equal 1 physical core.

- Replacing a node:  
On the **Nodes** tab, locate the node to be replaced. In the **Operation** column, click **Replace**. No fee is charged for this operation.  
Check the node replacement records on the **Records** page. **Running** indicates that the node is being replaced. After the replacement, you can check the new node in the node list.  
The replacement can last no longer than 24 hours. If no suitable resource is found after the replacement times out, the status changes to **Failed**. Hover over  to check the failure cause.

 **NOTE**

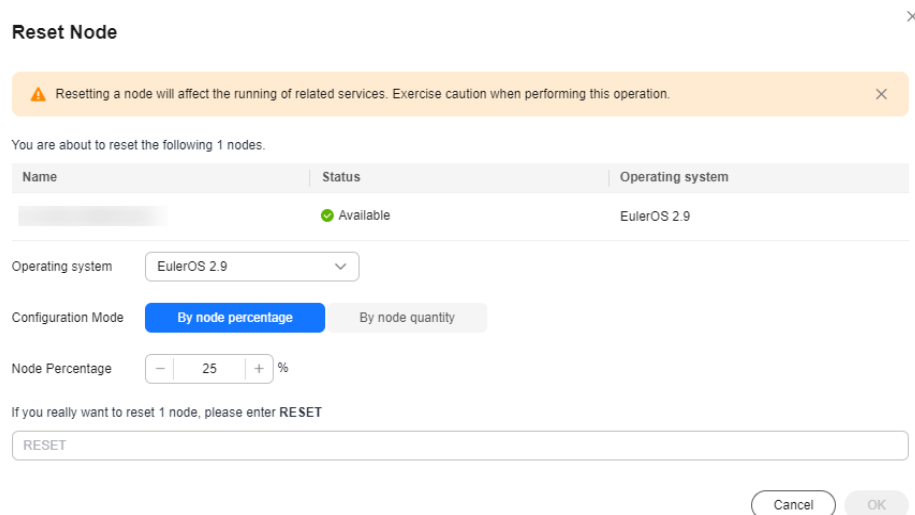
- The number of replacements per day cannot exceed 20% of the total nodes in the resource pool. The number of nodes to be replaced cannot exceed 5% of the total nodes in the resource pool.
  - Ensure that there are idle node resources. Otherwise, the replacement may fail.
  - If there are any nodes in the **Resetting** state in the operation records, nodes in the resource pool cannot be replaced.
- Resetting a node  
On the **Nodes** tab, locate the node you want to reset. Click **Reset** in the **Operation** column to reset a node. You can also select multiple nodes, and click **Reset** to reset multiple nodes.  
Configure the parameters described in the table below.

**Table 2-2** Parameters

| Parameter          | Description   |
|--------------------|---|
| Operating System   | Select an OS from the drop-down list box.   |
| Configuration Mode | Select a configuration mode for resetting the node. <ul style="list-style-type: none"> <li>• <b>By node percentage:</b> the maximum ratio of nodes that can be reset if there are multiple nodes in the reset task</li> <li>• <b>By node quantity:</b> the maximum number of nodes that can be reset if there are multiple nodes in the reset task</li> </ul> |

Check the node reset records on the **Records** page. If the node is being reset, its status is **Resetting**. After the reset is complete, the node status changes to **Available**. Resetting a node will not be charged.

Figure 2-8 Resetting a node



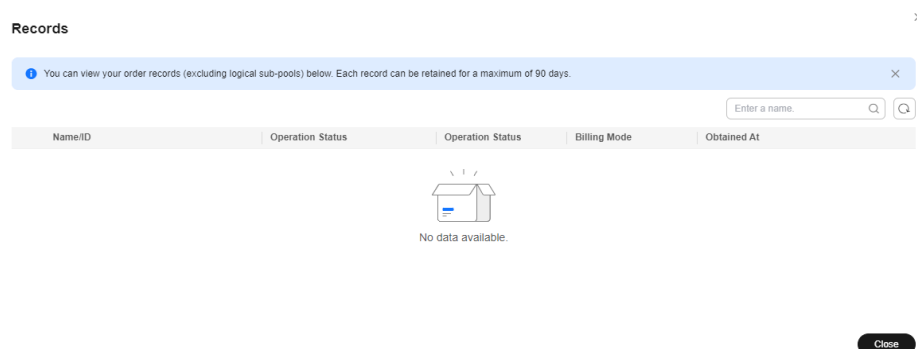
**NOTE**

- Resetting a node will affect running services.
- Only nodes in the **Available** state can be reset.
- A single node can be in only one reset task at a time. Multiple reset tasks cannot be delivered to the same node at a time.
- If there are any nodes in the **Replacing** state in the operation records, nodes in the resource pool cannot be reset.
- When the driver of a resource pool is being upgraded, nodes in this resource pool cannot be reset.
- For GPU and NPU specifications, after the node is reset, the driver of the node may be upgraded. Wait patiently.

Figure 2-9 Nodes



Figure 2-10 Operation records



- Deleting, unsubscribing from, or releasing a node
  - For a pay-per-use resource pool, click **Delete** in the **Operation** column.



To delete nodes in batches, select the check boxes next to the node names, and click **Delete**.

- For a yearly/monthly resource pool whose resources are not expired, click **Unsubscribe** in the **Operation** column.
- For a yearly/monthly resource pool whose resources are expired (in the grace period), click **Release** in the **Operation** column.

If the delete button is available for a yearly/monthly node, the node is an inventory node, click **Delete**.


**NOTE**

- Before deleting, unsubscribing from, or releasing a node, ensure that there are no running jobs on this node. Otherwise, the jobs will be interrupted.
- Delete, unsubscribe from, or release abnormal nodes in a resource pool and add new ones for substitution.
- If there is only one node, it cannot be deleted, unsubscribed from, or released.

## Viewing Resource Pool Specifications

On the resource pool details page, click **Specifications**. You can view the specifications used by the resource pool and the number of each specification.

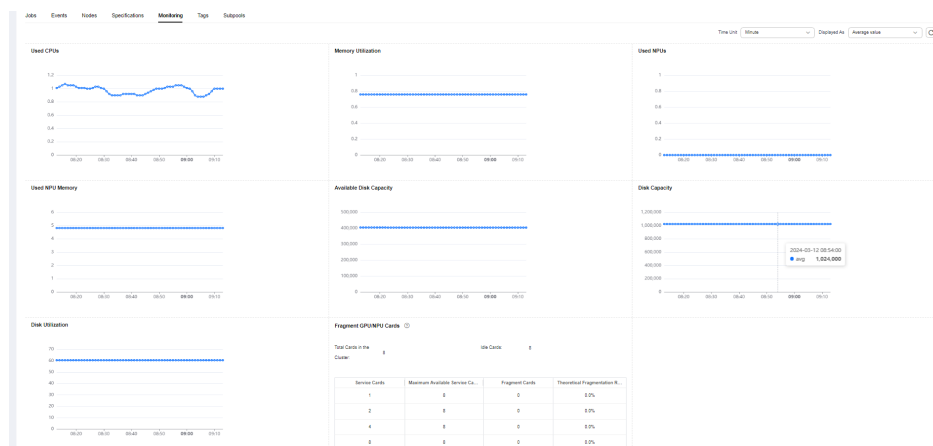
**Figure 2-11** View resource pool specifications (The container engine size is displayed as the default value if it is not set.)



## Viewing Resource Pool Monitoring Information

On the resource pool details page, click **Monitoring**. The resource usage including used CPUs, memory usage, and available disk capacity of the resource pool is displayed. If AI accelerators are used in the resource pool, the GPU and NPU monitoring information is also displayed.

**Figure 2-12** Viewing resource views




## Viewing Tags

You can add tags to a resource pool for quick search.

On the resource pool details page, click **Tags**. You can view, add, modify, and delete tags of a resource pool. For details about how to use tags, see [How Does ModelArts Use Tags to Manage Resources by Group?](#)

**Figure 2-13** Tags



### NOTE

You can add up to 20 tags.

## 2.4 Resizing a Resource Pool

### Description

The demand for resources in a dedicated resource pool may change due to the changes of AI development services. In this case, you can resize your dedicated resource pool in ModelArts.

- You can add nodes for existing flavors in the resource pool.
- You can delete nodes for existing flavors in the resource pool.

### NOTE

Before scaling in a resource pool, ensure that there are no services running in the pool. Alternatively, go to the resource pool details page, delete the nodes where no services are running to scale in the pool.

### Constraints

- Only dedicated resource pools in the **Running** status can be resized.
- When scaling in a dedicated resource pool, the number of flavors or nodes of a flavor cannot be decreased to 0.

### Resizing a Dedicated Resource Pool


You can resize a resource pool in any of the following ways:

- Adjusting the number of nodes of existing specifications
  - Resizing the container engine space
1. Log in to the ModelArts management console. In the navigation pane, choose **Dedicated Resource Pools > Elastic Cluster**.

### NOTE

A resource pool is suspended when it is migrated from the old version to the new version. You cannot adjust the capacity of such a resource pool or unsubscribe from it.

**Figure 2-14** Resource Pools




2. Add or delete nodes.

Click **Adjust Capacity** in the **Operation** column of the target resource pool.

In the **Resource Configurations** area, set **AZ** to **Automatically allocated** or **Specifies AZ**. Click **Submit** and then **OK** to save the changes.

- If **AZ** is set to **Automatically allocated**, you can increase or decrease the number of nodes to scale out or in the resource pool. After the scaling, nodes are automatically allocated to AZs.
- If you select **Specifies AZ**, you can allocate nodes to different AZs.

**Figure 2-15** Resource Configurations



3. Resize the container engine space.

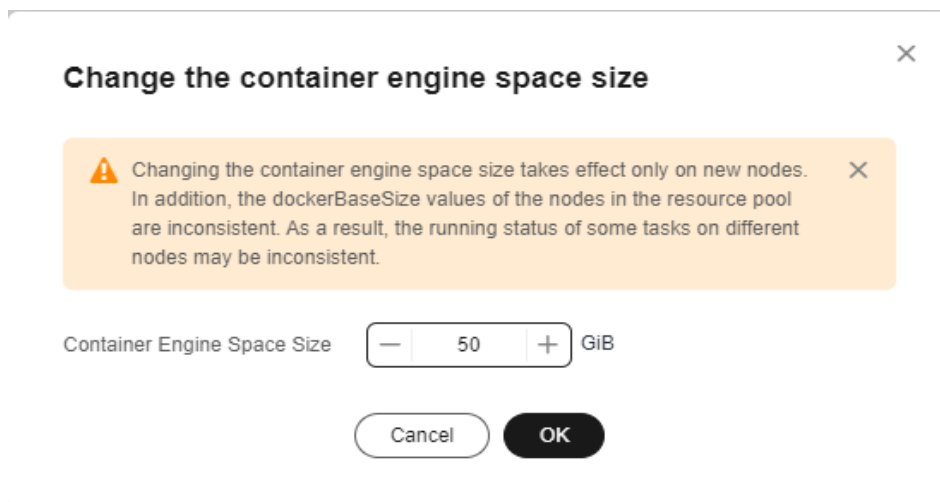
If you need larger container engine size, perform any of the following operations:

- For new resources, you can specify the container engine space when creating a resource pool. For details, see advanced configurations of **Specification Management** in [Creating a Resource Pool](#).
- For existing resources, the container engine space can be modified.
  - Method 1: Click the target resource pool to view its details. Click the **Specifications** tab, locate the target specifications, and click **Change the container engine space size** in the **Operation** column.
  - Method 2: Locate the target resource pool and click **Adjust Capacity** in the **Operation** column.

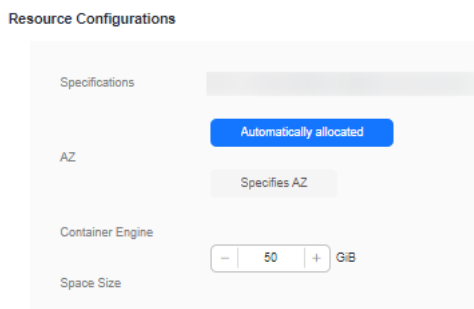
**NOTICE**

Resizing the container engine space is only applicable to new nodes. Furthermore, dockerBaseSize may vary across nodes of this flavor within the resource pool. Consequently, this can lead to discrepancies in the status of tasks distributed among different nodes.

**Figure 2-16** Resizing the container engine space (**Specifications** tab)



**Figure 2-17** Resizing the container engine space (**Resize** page)



## 2.5 Setting a Renewal Policy

### Description

ModelArts allows you to perform the following operations for yearly/monthly resource pools:

- Enable auto-renewal.
- Modify auto-renewal settings.
- Manually renew them.

### Constraints

The target dedicated resource pool must be running.

### Procedure

1. Log in to the ModelArts console. In the navigation pane, choose **Dedicated Resource Pools > Elastic Cluster**.
2. In the resource pool list, choose **More > Set Renewal Policy** in the **Operation** column of the target resource pool.

3. In the dialog box that appears, click **OK**. You will see the **Renewals** page of the billing center.
4. Set the renewal policy.
  - To enable auto-renewal for a yearly/monthly resource pool, click the **Manual Renewals** tab, locate the target resource pool, and choose **More > Enable Auto-Renewal** in the **Operation** column.
  - To modify auto-renewal settings for a yearly/monthly resource pool, click the **Auto Renewals** tab, locate the target resource pool, and choose **More > Modify Auto-Renew** in the **Operation** column to modify auto-renewal settings, such as the renewal mode, renewal duration, and number of renewals.
  - To manually renew a yearly/monthly resource pool, locate it and click **Renew** in the **Operation** column.

## 2.6 Modifying the Expiration Policy

### Description

ModelArts allows you to change the expiration policy of a yearly/monthly resource pool to pay-per-use or non-renewal after expiration.

### Constraints

The target dedicated resource pool must be running.

### Procedure

1. Log in to the ModelArts console. In the navigation pane, choose **Dedicated Resource Pools > Elastic Cluster**.
2. In the resource pool list, choose **More > Change Billing Mode** in the **Operation** column of the target resource pool.
3. In the dialog box that appears, click **OK**. You will see the **Renewals** page of the billing center.
4. Modify the expiration policy.
  - If auto-renewal has not been enabled for the target resource pool, click the **Manual Renewals** tab, and choose **More > Change to Pay-per-Use After Expiration** or **More > Cancel Renewal** in the **Operation** column of the target resource pool.
  - If auto-renewal has been enabled for the target resource pool, click the **Auto Renewals** tab, and choose **More > Cancel Renewal** in the **Operation** column of the target resource pool.

## 2.7 Migrating the Workspace

### Context

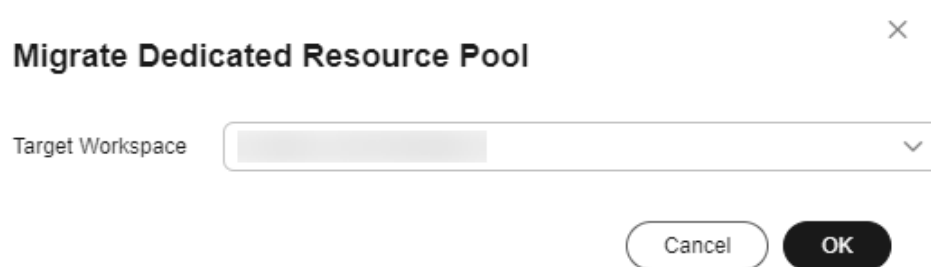
The workspace of a dedicated resource pool is associated with an enterprise project, which involves bill collection. ModelArts provides workspaces to isolate

resource operation permissions of different IAM users. Workspace migration includes resource pool migration and network migration. For details, see the following sections.

## Migrating the Workspace for a Resource Pool

1. Log in to the ModelArts management console. In the navigation pane, choose **Dedicated Resource Pools > Elastic Cluster**.
2. In the resource pool list, choose **More > Migrate Workspace** in the **Operation** column of the target resource pool.
3. In the **Migrate Dedicated Resource Pool** dialog box that appears, select the target workspace and click **OK**.

Figure 2-18 Migrating the workspace



## Migrating the Workspace for a Network

1. Log in to the ModelArts management console. In the navigation pane, choose **Dedicated Resource Pools > Elastic Cluster**. Then, click the **Network** tab.
2. In the network list, choose **More > Migrate Workspace** in the **Operation** column of the target network.
3. In the dialog box that appears, select the target workspace and click **OK**.

Figure 2-19 Migrating the workspace



## 2.8 Changing Job Types Supported by a Resource Pool

### Description

ModelArts supports many types of jobs. Some of them can run in dedicated resource pools, including training jobs, inference services, and notebook development environments.

You can change job types supported by a dedicated resource pool. Available options for **Job Type** are **Training Job**, **Inference Service**, and **DevEnviron**.

Only selected types of jobs can be delivered to the corresponding dedicated resource pool.

---

 **CAUTION**

To support different job types, different operations are performed in the backend, such as installing plug-ins and setting the network environment. Some operations use resources in the resource pool. As a result, available resources for you decrease. Therefore, select only the job types you need to avoid resource waste.

---

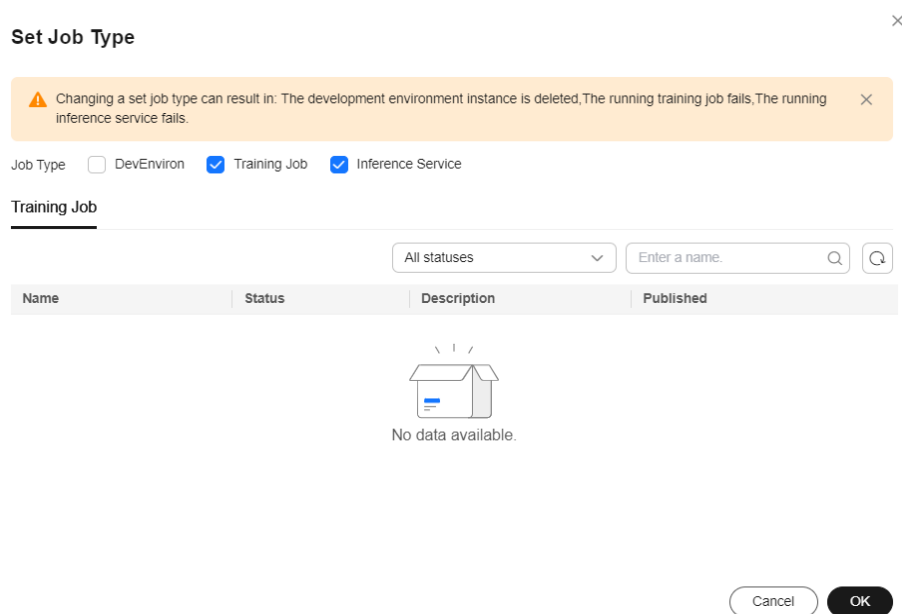
### Constraints

The target dedicated resource pool must be running.

### Procedure

1. Log in to the ModelArts management console. In the navigation pane, choose **Dedicated Resource Pools > Elastic Cluster**.
2. In the **Operation** column of a resource pool, choose **More > Set Job Type**.
3. In the **Set Job Type** dialog box, select job types.

**Figure 2-20** Setting the job type



4. Click **OK**.

## 2.9 Upgrading a Resource Pool Driver

### Description

If GPUs or Ascend resources are used in a dedicated resource pool, you may need to customize GPU or Ascend drivers. ModelArts allows you to upgrade GPU or Ascend drivers of your dedicated resource pools.

There are two driver upgrade modes: secure upgrade and forcible upgrade.

#### NOTE

- Secure upgrade: Running services are not affected. After the upgrade starts, the nodes are isolated (new jobs cannot be delivered). After the existing jobs on the nodes are complete, the upgrade is performed. The secure upgrade may take a long time because the jobs must be completed first.
- Forcible upgrade: The drivers are directly upgraded, regardless of whether there are running jobs.

### Constraints

- The target dedicated resource pool must be running, and the resource pool contains GPU or Ascend resources.
- For a logical resource pool, the driver can be upgraded only after node binding is enabled. To enable node binding, submit a service ticket to contact Huawei engineers.

### Upgrading the Driver

1. Log in to the ModelArts management console. In the navigation pane, choose **Dedicated Resource Pools > Elastic Cluster**.



2. In the **Operation** column of the target resource pool, choose **More > Upgrade Driver**.
3. In the **Upgrade Driver** dialog box, the driver type, number of nodes, current version, target version, and upgrade mode of the dedicated resource pool are displayed.
  - **Target Version:** Select a target driver version from the drop-down list.
  - **Upgrade Mode:** Select **Secure upgrade** or **Forcible upgrade**.
  - **Rolling Mode:** Once enabled, you can upgrade the driver in rolling mode. Currently, rolling by node percentage and by node quantity are supported. If **By node percentage** is selected, the number of nodes to be upgraded in each batch is the node ratio multiplied by the total number of nodes in the resource pool. If **By node quantity** is selected, the number of nodes to be upgraded in each batch is what you configured.

Figure 2-21 Upgrading a driver

Upgrade Driver - [blurred] X

Driver Type GPU

Nodes 1

Current Version [blurred]

Target Version [dropdown]

Upgrade Mode Secure upgrade Forcible upgrade ?

Rolling

Rolling Mode By node percentage By node quantity

Node Percentage - 25 + %

Cancel OK

4. Click **OK** to start the driver upgrade.

## 2.10 Deleting a Resource Pool

If a dedicated resource pool is no longer needed for AI service development, you can delete the resource pool to release resources.

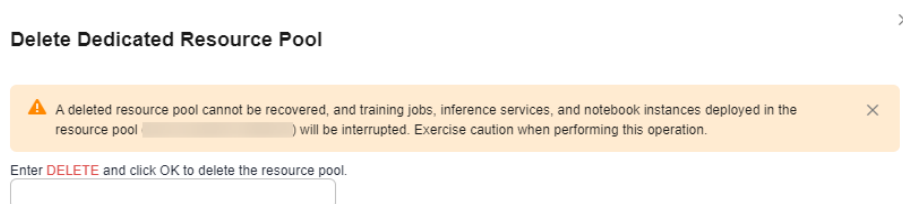
### NOTE

After a dedicated resource pool is deleted, the development environments, training jobs, and inference services that depend on the resource pool are unavailable. A dedicated resource pool cannot be restored after being deleted.

1. Log in to the ModelArts management console. In the navigation pane, choose **Dedicated Resource Pools > Elastic Cluster**.
2. Locate the row that contains the target resource pool, choose **More > Delete** in the **Operation** column.
3. In the **Delete Dedicated Resource Pool** dialog box, enter **DELETE** in the text box and click **OK**.

You can switch between tabs on the details page to view the training jobs and notebook instances created using the resource pool, and inference services deployed in the resource pool.

**Figure 2-22** Deleting a resource pool



## 2.11 Abnormal Status of a Dedicated Resource Pool

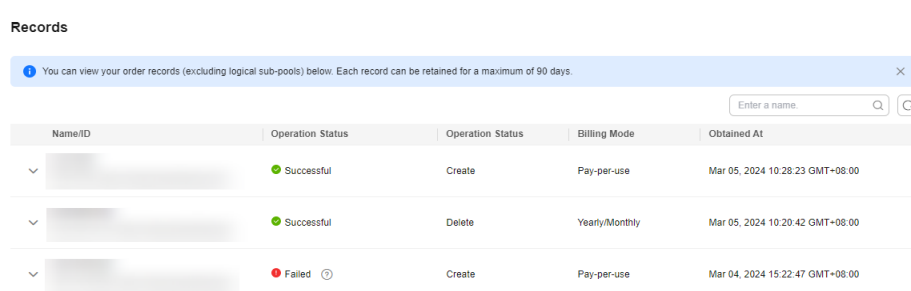
### Resource Quota Limit

When you use a dedicated resource pool (for example, scaling resources, creating a VPC, creating a VPC and subnet, or interconnecting a VPC), if the system displays a message indicating that the resource quota is limited, [submit a service ticket](#).

### Creation Failed/Change Failed

1. Log in to the ModelArts management console. In the navigation pane, choose **Dedicated Resource Pools > Elastic Cluster**.
2. Click **Records** on the right of **Create**. On the **Records** dialog box, view failed task records.

**Figure 2-23** Creating a resource pool failed



3. Hover the cursor over , view the cause of task failures.

 **NOTE**

By default, failed task records are sorted by application time. A maximum of 500 failed task records can be displayed and retained for three days.

## Locating Faulty Node

ModelArts will add a taint on a detected K8S faulty node so that jobs will not be affected or scheduled to the tainted node. The following table lists the faults can be detected. You can locate the fault by referring to the isolation code and detection method.

**Table 2-3** Isolation code

| Isolation Code | Category | Sub-Category | Description           | Detection Method  |
|----------------|----------|--------------|-----------------------|---|
| A050101        | GPU      | GPU memory   | GPU ECC error exists. | <p>Run the <b>nvidia-smi -a</b> command and check whether <b>Pending Page Blacklist</b> is <b>Yes</b> or the value of <b>multi-bit Register File</b> is greater than 0. For Ampere GPUs, check whether the following content exists:</p> <ul style="list-style-type: none"> <li>• Uncorrectable SRAM error</li> <li>• Remapping Failure records</li> <li>• Xid 95 events in <b>dmsg</b></li> </ul> <p>(For details, see <a href="#">NVIDIA GPU Memory Error Management</a>.)</p> <p>The Ampere architecture has the following levels of GPU memory errors:</p> <ul style="list-style-type: none"> <li>• L1: These are single-bit ECC errors that can be corrected. They do not affect the running services. To check for these errors, run the <b>nvidia-smi -a</b> command and look for <b>Volatile Correctable</b>.</li> <li>• L2: These are multi-bit ECC errors that cannot be corrected. They cause the running services to fail and require a process restart to recover. To check for these errors, run the <b>nvidia-smi -a</b> command and look for <b>Volatile Uncorrectable</b>.</li> <li>• L3: These are unsuppressed errors and may affect other services. They require a card reset or a node reboot to clear. To check for these errors, look for the Xid events that contain the number 95. (The Remapped Pending records are only for reference. You need to reset the cards when the service is idle to trigger the remapping process.)</li> <li>• L4: These are errors that require a card replacement. To check for these errors, look for the <b>SRAM</b></li> </ul> |

| Isolation Code | Category | Sub-Category | Description   | Detection Method   |
|----------------|----------|--------------|---|--|
|                |          |              |   | <b>Uncorrectable</b> field that is greater than 4 or the <b>Remapped Failed</b> field that is not zero.  |
| A05 0102       | GPU      | Other        | The <b>nvidia-smi</b> output contains ERR.                      | Run <b>nvidia-smi -a</b> and check whether the output contains ERR. Normally, the hardware, such as the power supply or the fan, is faulty.                  |
| A05 0103       | GPU      | Other        | The execution of <b>nvidia-smi</b> times out or does not exist. | Check that exit code of <b>nvidia-smi</b> is not <b>0</b> .  |
| A05 0104       | GPU      | GPU Memory   | ECC error occurred 64 times.                                    | Run the <b>nvidia-smi -a</b> command, locate <b>Retired Pages</b> , and check whether the sum of <b>Single Bit</b> and <b>Double Bit</b> is greater than 64. |
| A05 0148       | GPU      | Other        | An infoROM alarm occurs.  | Run the <b>nvidia-smi</b> command and check whether the output contains the alarm "infoROM is corrupted".  |
| A05 0109       | GPU      | Other        | Other GPU errors  | Check whether other GPU error exists. Normally, there is a faulty hardware. Contact the technical engineer.  |
| A05 0147       | IB       | Link         | The IB NIC is abnormal.   | Run the <b>ibstat</b> command and check whether the NIC is not in active state.  |
| A05 0121       | NPU      | Other        | A driver exception is detected by NPU DCMI.                     | The NPU driver environment is abnormal.  |
| A05 0122       | NPU      | Other        | The NPU DCMI device is abnormal.                                | The NPU device is abnormal. The Ascend DCMI interface returns a major or urgent alarm.   |
| A05 0123       | NPU      | Link         | The NPU DCMI net is abnormal.                                   | The NPU network connection is abnormal.  |
| A05 0129       | NPU      | Other        | Other NPU errors  | Check whether other NPU error exists. You cannot rectify the fault. Contact the technical engineer.  |

| Isolation Code | Category        | Sub-Category   | Description   | Detection Method   |
|----------------|-----------------|----------------|---|--|
| A050149        | NPU             | Link           | Check whether the network port of the hccn tool is intermittently disconnected. | The NPU network is unstable and intermittently disconnected. Run the <b>hccn_tool-i \${device_id} -link_stat -g</b> command and the network is disconnected more than five times within 24 hours.  |
| A050951        | NPU             | GPU memory     | The number of NPU ECCs reaches the maintenance threshold.                       | The NPU's HBM Double Bit Isolated Pages Count value is greater than or equal to 64.  |
| A050146        | Runtime         | Other          | The NTP is abnormal.  | The ntpd or chronyd service is abnormal.   |
| A050202        | Runtime         | Other          | The node is not ready.  | The node is unavailable. The K8S node contains one of the following taints: <ul style="list-style-type: none"> <li>node.kubernetes.io/unreachable</li> <li>node.kubernetes.io/not-ready</li> </ul> |
| A050203        | Runtime         | Disconnection  | The number of normal AI cards does not match the actual capacity.               | The GPU or NPU is disconnected.  |
| A050206        | Runtime         | Other          | The Kubelet hard disk is read-only.   | The <b>/mnt/paas/kubernetes/kubelet</b> directory is read-only.  |
| A050801        | Node management | Node O&M       | Resource is reserved.   | The node is marked as the standby node and contains a taint.   |
| A050802        | Node management | Node O&M       | An unknown error occurs.  | The node is marked with an unknown taint.  |
| A200001        | Node management | Driver upgrade | The GPU is being upgraded.  | The GPU is being upgraded.   |

| Isolation Code | Category         | Sub-Category             | Description  | Detection Method   |
|----------------|------------------|--------------------------|--|--|
| A200002        | Node management  | Driver upgrade           | The NPU is being upgraded.                                 | The NPU is being upgraded.   |
| A200008        | Node management  | Node admission           | The admission is being examined.                           | The admission is being examined, including basic node configuration check and simple service verification. |
| A050933        | Node management  | Fault tolerance Failover | The Failover service on the tainted node will be migrated. | The Failover service on the tainted node will be migrated.   |
| A050931        | Training toolkit | Pre-check container      | A GPU error is detected in the pre-check container.        | A GPU error is detected in the pre-check container.  |
| A050932        | Training toolkit | Pre-check container      | An IB error is detected in the pre-check container.        | An IB error is detected in the pre-check container.  |

## 2.12 ModelArts Network

### ModelArts Network and VPC

ModelArts networks are backed by VPCs and used for interconnecting nodes in a ModelArts resource pool. You can only configure the name and CIDR block for a network. To ensure that there is no IP address segment in the CIDR block overlapped with that of the VPC to be accessed, multiple CIDR blocks are available for you to select.

A VPC provides a logically isolated virtual network for your instances. You can configure and manage the network as required. VPC provides logically isolated, configurable, and manageable virtual networks for cloud servers, cloud containers, and cloud databases. It helps you improve cloud service security and simplify network deployment.

### Prerequisites


- A VPC is available.

- A subnet is available.

## Creating a Network

1. Log in to the ModelArts management console. In the navigation pane, choose **Dedicated Resource Pools > Elastic Cluster**.
2. Click **Network** and then **Create**.

Figure 2-24 Network list



3. In the **Create Network** dialog box, set parameters.
  - **Network Name:** customizable name
  - **CIDR Block:** You can select **Preset** or **Custom**.

### NOTE


- Each user can create a maximum of 15 networks.
  - Ensure there is no IP address segment in the CIDR block overlaps that of the VPC to be accessed. The CIDR block cannot be changed after the network is created. Possible conflict CIDR blocks are as follows:
    - Your VPC CIDR block
    - Container CIDR block (consistently to be 172.16.0.0/16)
    - Service CIDR block (consistently to be 10.247.0.0/16)
4. Confirm the settings and click **OK**.

## (Optional) Interconnecting a VPC with a ModelArts Network

VPC interconnection allows you to use resources across VPCs, improving resource utilization.

1. On the **Network** page, click **Interconnect VPC** in the **Operation** column of the target network.

Figure 2-25 Interconnect VPC



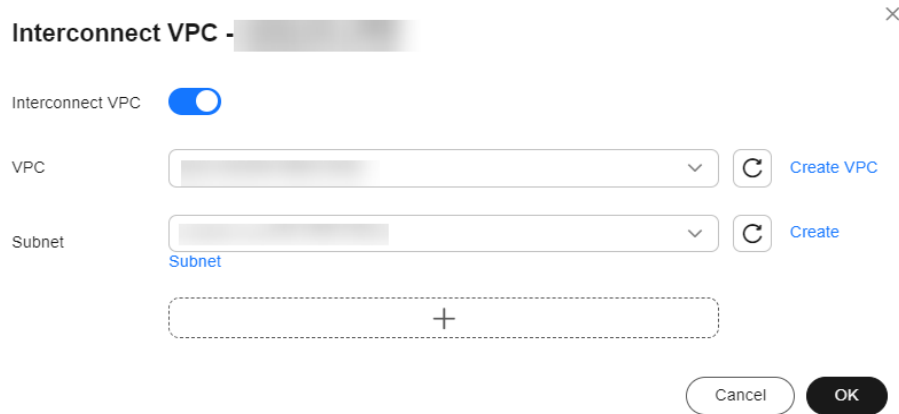
2. In the displayed dialog box, click the button on the right of **Interconnect VPC**, and select an available VPC and subnet from the drop-down lists.

### NOTE

The peer network to be interconnected cannot overlap with the current CIDR block.



**Figure 2-26** Parameters for interconnecting a VPC with a network



- If no VPC is available, click **Create VPC** on the right to create a VPC.
- If no subnet is available, click **Create Subnet** on the right to create a subnet.
- Multiple subnets in a VPC can be interconnected. You can click + to add up to 10 subnets.

## Enabling a Dedicated Resource Pool to Access the Internet

To enable a dedicated resource pool to access the Internet, follow these steps:

- Step 1** Interconnect a VPC. For details, see [\(Optional\) Interconnecting a VPC with a ModelArts Network](#).
- Step 2** For details about how to configure an SNAT server for a VPC, see [Configuring an SNAT Server](#).

----End

## Deleting a Network


If a network is no longer needed for AI service development, you can delete the network.

1. Go to the **Network** tab page and click **Delete** in the **Operation** column of a network.
2. Confirm the information and click **OK**.

## 2.13 ModelArts Nodes

Nodes that are not managed by the resource pool are considered as free nodes. To view the information about free nodes, log in to the ModelArts management console, choose **Dedicated Resource Pools > Elastic Cluster**, and click the **Nodes** tab.

**Figure 2-27** Nodes



Release the free nodes resources according to the following content:

- For a pay-per-use node, click **Delete** in the **Operation** column.
- For a yearly/monthly node whose resources are not expired, click **Unsubscribe** in the **Operation** column.
- For a yearly/monthly node whose resources are expired (in the grace period), click **Release** in the **Operation** column.

If the delete button is available for a yearly/monthly node, click the button to delete the node.

 **NOTE**

Deletion, unsubscription, and release operations cannot be undone. Exercise caution when performing this operation.

# 3 Audit Logs

## 3.1 Key Operations Recorded by CTS

With CTS, you can obtain operations associated with ModelArts for later query, audit, and backtrack operations.

### Prerequisites

CTS has been enabled.

### Key Data Management Operations Traced by CTS

**Table 3-1** Key data management operations traced by CTS

| Operation                      | Resource Type | Trace                  |
|--------------------------------|---------------|------------------------|
| Creating a dataset             | Dataset       | createDataset          |
| Deleting a dataset             | Dataset       | deleteDataset          |
| Updating a dataset             | Dataset       | updateDataset          |
| Publishing a dataset version   | Dataset       | publishDatasetVersion  |
| Deleting a dataset version     | Dataset       | deleteDatasetVersion   |
| Synchronizing the data source  | Dataset       | syncDataSource         |
| Exporting a dataset            | Dataset       | exportDataFromDataset  |
| Creating an auto labeling task | Dataset       | createAutoLabelingTask |
| Creating an auto grouping task | Dataset       | createAutoGroupingTask |

| Operation   | Resource Type | Trace                        |
|---|---------------|------------------------------|
| Creating an auto deployment task  | Dataset       | createAutoDeployTask         |
| Importing samples to a dataset  | Dataset       | importSamplesToDataset       |
| Creating a dataset label  | Dataset       | createLabel                  |
| Updating a dataset label  | Dataset       | updateLabel                  |
| Deleting a dataset label  | Dataset       | deleteLabel                  |
| Deleting a dataset label and their labeled samples                        | Dataset       | deleteLabelWithSamples       |
| Adding samples  | Dataset       | uploadSamples                |
| Deleting samples  | Dataset       | deleteSamples                |
| Stopping an auto labeling task  | Dataset       | stopTask                     |
| Creating a team labeling task   | Dataset       | createWorkforceTask          |
| Deleting a team labeling task   | Dataset       | deleteWorkforceTask          |
| Starting the acceptance of a team labeling task                           | Dataset       | startWorkforceSampling-Task  |
| Approving, rejecting, or canceling the acceptance of a team labeling task | Dataset       | updateWorkforceSam-plingTask |
| Submitting sample review comments for an acceptance task                  | Dataset       | acceptSamples                |
| Adding a label to a sample  | Dataset       | updateSamples                |
| Sending an email to team labeling members                                 | Dataset       | sendEmails                   |
| Starting a team labeling task as the team manager                         | Dataset       | startWorkforceTask           |
| Updating a team labeling task   | Dataset       | updateWorkforceTask          |
| Adding a label to a team-labeled sample                                   | Dataset       | updateWorkforceTask-Samples  |

| Operation   | Resource Type      | Trace                    |
|---|--------------------|--------------------------|
| Reviewing team labeling results   | Dataset            | reviewSamples            |
| Creating a labeling team member   | Workforce          | createWorker             |
| Updating labeling team members  | Workforce          | updateWorker             |
| Deleting a labeling team member   | Workforce          | deleteWorker             |
| Deleting labeling team members in a batch   | Workforce          | batchDeleteWorker        |
| Creating a labeling team  | Workforce          | createWorkforce          |
| Updating a labeling team  | Workforce          | updateWorkforce          |
| Deleting a labeling team  | Workforce          | deleteWorkforce          |
| Automatically creating an IAM agency  | IAM                | createAgency             |
| Logging in to the labeling console as a team labeling member  | labelConsoleWorker | workerLoginLabelConsole  |
| Logging out of the labeling console as a team labeling member   | labelConsoleWorker | workerLogoutLabelConsole |
| Changing the password for logging in to the labeling console as a team labeling member                        | labelConsoleWorker | workerChangePassword     |
| Handling the issue that the password for logging in to the labeling console as a team labeling member is lost | labelConsoleWorker | workerForgetPassword     |
| Resetting the password for logging in to the labeling console through the URL as a team labeling member       | labelConsoleWorker | workerResetPassword      |

## Key DevEnviron Operations Traced by CTS

**Table 3-2** Key DevEnviron operations traced by CTS

| Operation                        | Resource Type | Trace             |
|----------------------------------|---------------|-------------------|
| Creating a notebook instance     | Notebook      | createNotebook    |
| Deleting a notebook instance     | Notebook      | deleteNotebook    |
| Opening a notebook instance      | Notebook      | openNotebook      |
| Starting a notebook instance     | Notebook      | startNotebook     |
| Stopping a notebook instance     | Notebook      | stopNotebook      |
| Updating a notebook instance     | Notebook      | updateNotebook    |
| Deleting a NotebookApp           | NotebookApp   | deleteNotebookApp |
| Switching CodeLab specifications | NotebookApp   | updateNotebookApp |

## Key Training Job Operations Traced by CTS

**Table 3-3** Key training job operations traced by CTS

| Operation                                   | Resource Type        | Trace                       |
|---|----------------------|-----------------------------|
| Creating a training job                     | ModelArtsTrainJob    | createModelArtsTrainJob     |
| Creating a training job version             | ModelArtsTrainJob    | createModelArtsTrainVersion |
| Stopping a training job                     | ModelArtsTrainJob    | stopModelArtsTrainVersion   |
| Modifying the description of a training job | ModelArtsTrainJob    | updateModelArtsTrainDesc    |
| Deleting a training job version             | ModelArtsTrainJob    | deleteModelArtsTrainVersion |
| Deleting a training job                     | ModelArtsTrainJob    | deleteModelArtsTrainJob     |
| Configuring training job                    | ModelArtsTrainConfig | createModelArtsTrainConfig  |

| Operation  | Resource Type            | Trace                          |
|--|--------------------------|--------------------------------|
| Modifying a training job configuration           | ModelArtsTrainConfig     | updateModelArtsTrain-Config    |
| Deleting a training job configuration            | ModelArtsTrainConfig     | deleteModelArtsTrain-Config    |
| Creating a visualization job                     | ModelArtsTensorboard-Job | createModelArtsTensorboardJob  |
| Deleting a visualization job                     | ModelArtsTensorboard-Job | deleteModelArtsTensorboardJob  |
| Modifying the description of a visualization job | ModelArtsTensorboard-Job | updateModelArtsTensorboardDesc |
| Stopping a visualization job                     | ModelArtsTensorboard-Job | stopModelArtsTensorboardJob    |
| Restarting a visualization job                   | ModelArtsTensorboard-Job | restartModelArtsTensorboardJob |

## Key AI Application Management Operations Traced by CTS

**Table 3-4** Key AI application management operations traced by CTS

| Operation                        | Resource Type | Trace         |
|----------------------------------|---------------|---------------|
| Creating an AI application       | Model         | addModel      |
| Updating an AI application       | Model         | updateModel   |
| Deleting an AI application       | Model         | deleteModel   |
| Creating a model conversion task | Convert       | addConvert    |
| Updating a model conversion task | Convert       | updateConvert |
| Deleting a model conversion task | Convert       | deleteConvert |

## Key Service Management Operations Traced by CTS

**Table 3-5** Key service management operations traced by CTS

| Operation  | Resource Type | Trace               |
|--|---------------|---------------------|
| Deploying a service                                      | Service       | addService          |
| Deleting a service                                       | Service       | deleteService       |
| Updating a service                                       | Service       | updateService       |
| Starting or stopping a service                           | Service       | startOrStopService  |
| Adding a user access key                                 | Service       | addAkSk             |
| Deleting a user access key                               | Service       | deleteAkSk          |
| Creating a dedicated resource pool                       | Cluster       | createCluster       |
| Deleting a dedicated resource pool                       | Cluster       | deleteCluster       |
| Adding a node to a dedicated resource pool               | Cluster       | addClusterNode      |
| Deleting a node from a dedicated resource pool           | Cluster       | deleteClusterNode   |
| Obtaining a result of creating a dedicated resource pool | Cluster       | createClusterResult |

## Key AI Gallery Operations Traced by CTS

**Table 3-6** Key AI Gallery operations traced by CTS

| Operation                        | Resource Type    | Trace                |
|----------------------------------|------------------|----------------------|
| Publishing an asset              | ModelArts_Market | create_content       |
| Modifying asset information      | ModelArts_Market | modify_content       |
| Publishing an asset version      | ModelArts_Market | add_version          |
| Subscribing to an asset          | ModelArts_Market | subscription_content |
| Removing an asset from favorites | ModelArts_Market | cancel_star_content  |



| Operation                  | Resource Type    | Trace               |
|----------------------------|------------------|---------------------|
| Liking an asset            | ModelArts_Market | like_content        |
| Unliking an asset          | ModelArts_Market | cancel_like_content |
| Publishing an activity     | ModelArts_Market | publish_activity    |
| Signing up an activity     | ModelArts_Market | regist_activity     |
| Modifying user information | ModelArts_Market | update_user         |

## Key Resource Management Operations Traced by CTS


**Table 3-7** Key resource management operations traced by CTS

| Operation                | Resource Type | Trace            |
|--------------------------|---------------|------------------|
| Creating a resource pool | PoolV2        | CreatePoolV2     |
| Deleting a resource pool | PoolV2        | DeletePoolV2     |
| Updating a resource pool | PoolV2        | UpdatePoolV2     |
| Creating a network       | NetworksV1    | CreateNetworksV1 |
| Deleting a network       | NetworksV1    | DeleteNetworksV1 |
| Update a network         | NetworksV1    | UpdateNetworksV1 |

## 3.2 Viewing Audit Logs


After CTS is enabled, CTS starts recording operations related to ModelArts. The CTS management console stores the last seven days of operation records. This section describes how to query operation records of the last seven days on the CTS management console.

### Procedure

1. Log in to the CTS management console.
2. Click  in the upper left corner of the page and select a region.
3. In the left navigation pane, click **Trace List**.
4. Specify the filter criteria used for querying traces. The following four filter criteria are available:
  - **Trace Source, Resource Type, and Search By**  
Select a filter criterion from the drop-down list.  
If you select **Trace name** for **Search By**, you need to select a specific trace name.

If you select **Resource ID** for **Search By**, you need to enter a specific resource ID.

If you select **Resource name** for **Search By**, you need to select or enter a specific resource name.

- **Operator**: Select a specific operator (a user rather than an account).
  - **Trace Status**: Available options include **All trace statuses**, **normal**, **warning**, and **incident**. You can only select one of them.
  - **Time Range**: You can view traces generated during any time range of the last seven days.
5. Click  on the left of a trace to expand its details.
  6. Click **View Trace** in the **Operation** column. In the displayed **View Trace** dialog box, the trace structure details are displayed.

For details about the key fields in the CTS trace structure, see [Cloud Trace Service User Guide](#).

# 4 Monitoring Resources

---

## 4.1 Overview

All metrics reported by ModelArts are stored in AOM, which enables you to consume metrics. You can view metric threshold alarms and reported alarms on the AOM console or use visualization tools such as Grafana to view and analyze the alarms. Grafana provides different views and templates for monitoring, which allow you to see the real-time resource usage on dashboards clearly.

## 4.2 Using Grafana to View AOM Monitoring Metrics

### 4.2.1 Procedure

Grafana supports various monitoring views and templates, meeting your diverse requirements. After adding the data source in Grafana, you can view all ModelArts monitoring metrics stored in AOM using Grafana.

To view AOM monitoring metrics using Grafana plugins, perform the following steps:

1. [Installing and Configuring Grafana](#)

 **NOTE**

You can install and configure Grafana using any of the following ways: [Installing and Configuring Grafana on Windows](#), [Installing and Configuring Grafana on Linux](#), and [Installing and Configuring Grafana on a Notebook Instance](#).

2. [Configuring a Grafana Data Source](#)
3. [Using Grafana to Configure Dashboards and View Metric Data](#)

### 4.2.2 Installing and Configuring Grafana

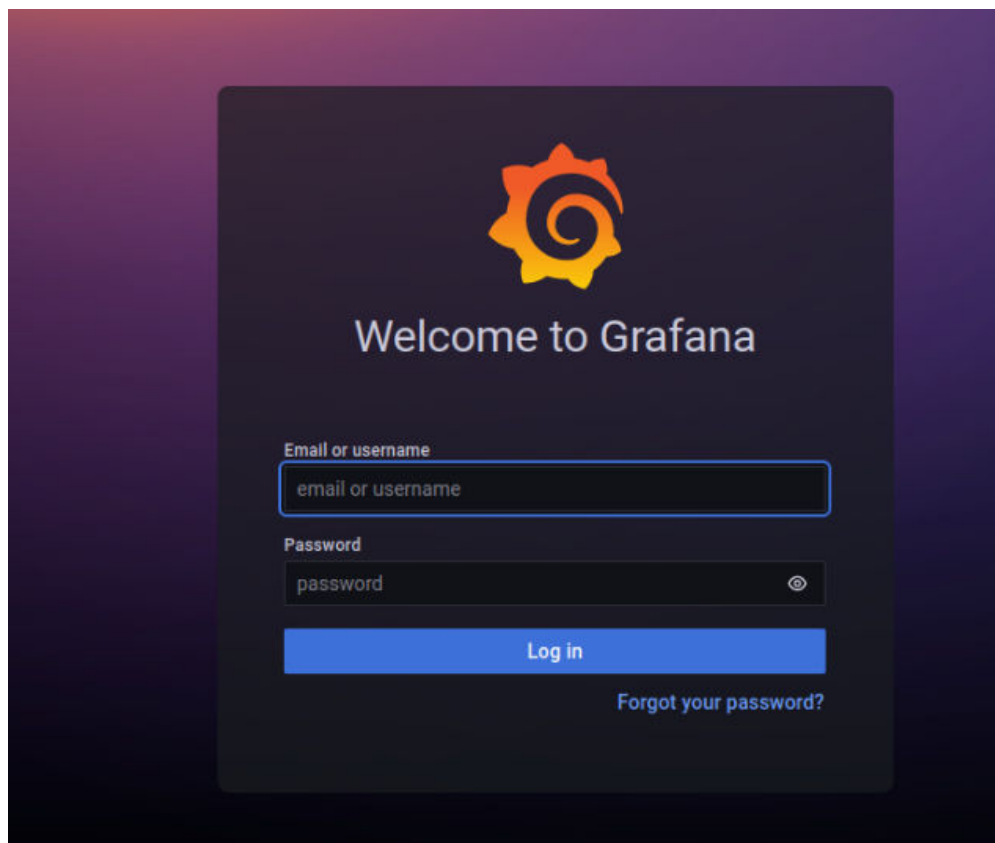
## 4.2.2.1 Installing and Configuring Grafana on Windows

### Application Scenario

This section describes how to install and configure Grafana on a Windows operating system.

### Procedure

1. Download the Grafana installation package.  
Go to the [download link](#), click **Download the installer**, and wait until the download is successful.
2. Install Grafana.  
Double-click the installation package and install Grafana as instructed.
3. In Windows Services Manager, enable Grafana.
4. Log in to Grafana.  
Grafana runs on port 3000 by default. After you open <http://localhost:3000>, the Grafana login page is displayed. The default username and password for the first login are **admin**. After the login is successful, change the password as prompted.



## 4.2.2.2 Installing and Configuring Grafana on Linux

### Prerequisites

- An Ubuntu server that is accessible to the Internet is available. If no, the following conditions must be met:
- You have obtained an ECS. (You are advised to select 8 vCPUs or higher, Ubuntu image of 22.04 version, and 100 GB local storage.) For details, see [Purchasing an ECS](#).
- You have purchased an EIP and bound it to the ECS. For details, see [Assigning an EIP and Binding It to an ECS](#).

### Procedure

1. Log in to the ECS. Select a login method. For details, see .
2. Run the following command to install libfontconfig1:

```
sudo apt-get install -y adduser libfontconfig1
```

The operation is successful if the following information is displayed:

```
root@ecs-9ec3:~# sudo apt-get install -y adduser libfontconfig1
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
adduser is already the newest version (3.118ubuntu5).
adduser set to manually installed.
libfontconfig1 is already the newest version (2.13.1-4.2ubuntu5).
libfontconfig1 set to manually installed.
The following packages were automatically installed and are no longer required:
  eatmydata libeatmydata1 libflashrom1 libftdi1-2 python-babel-localedata python3-babel python3-certifi python3-jinja2
  python3-json-pointer python3-jsonpatch python3-jsonschema python3-markupsafe python3-pyrsistent python3-requests python3-tz
  python3-urllib3
Use 'sudo apt autoremove' to remove them.
0 upgraded, 0 newly installed, 0 to remove and 4 not upgraded.
```

3. Run the following command to download the Grafana installation package:  
wget [https://dl.grafana.com/oss/release/grafana\\_9.3.6\\_amd64.deb](https://dl.grafana.com/oss/release/grafana_9.3.6_amd64.deb) --no-check-certificate

Download completed:

```
root@ecs-9ec3:~# wget https://dl.grafana.com/oss/release/grafana_9.3.6_amd64.deb --no-check-certificate
--2023-03-07 10:22:12-- https://dl.grafana.com/oss/release/grafana_9.3.6_amd64.deb
Resolving dl.grafana.com (dl.grafana.com)... 151.101.42.217
Connecting to dl.grafana.com (dl.grafana.com)|151.101.42.217|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 89252050 (85M) [application/octet-stream]
Saving to: 'grafana_9.3.6_amd64.deb'

grafana_9.3.6_amd64.deb 100%[=====] 85.12M 379KB/s in 2m 21s
2023-03-07 10:24:36 (617 KB/s) - 'grafana_9.3.6_amd64.deb' saved [89252050/89252050]
```

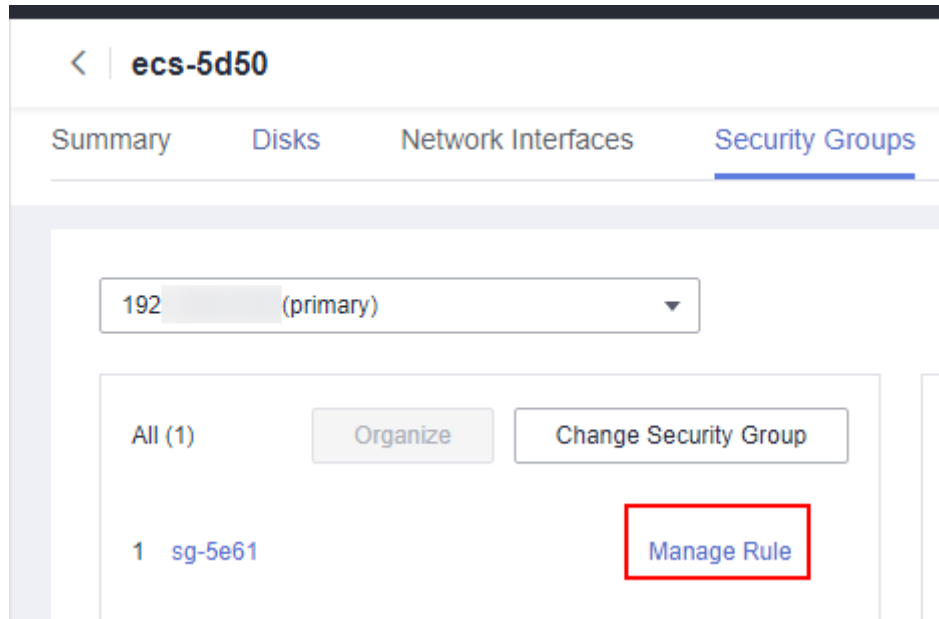
4. Run the following command to install Grafana:  
sudo dpkg -i grafana\_9.3.6\_amd64.deb

```
root@ecs-9ec3:~# sudo dpkg -i grafana_9.3.6_amd64.deb
Selecting previously unselected package grafana.
(Reading database ... 80788 files and directories currently installed.)
Preparing to unpack grafana_9.3.6_amd64.deb ...
Unpacking grafana (9.3.6) ...
Setting up grafana (9.3.6) ...
Adding system user `grafana' (UID 116) ...
Adding new user `grafana' (UID 116) with group `grafana' ...
Not creating home directory `/usr/share/grafana'.
### NOT starting on installation, please execute the following statements to configure grafana to start automatically using syst
emd
sudo /bin/systemctl daemon-reload
sudo /bin/systemctl enable grafana-server
### You can start grafana-server by executing
sudo /bin/systemctl start grafana-server
```


5. Run the following command to start Grafana:  
sudo /bin/systemctl start grafana-server
6. Access Grafana configurations on your local PC.

Ensure that an EIP has been bound to the ECS and the [security group](#) configuration is correct (the inbound traffic from TCP port 3000 and all outbound traffic are allowed). Configuration process:

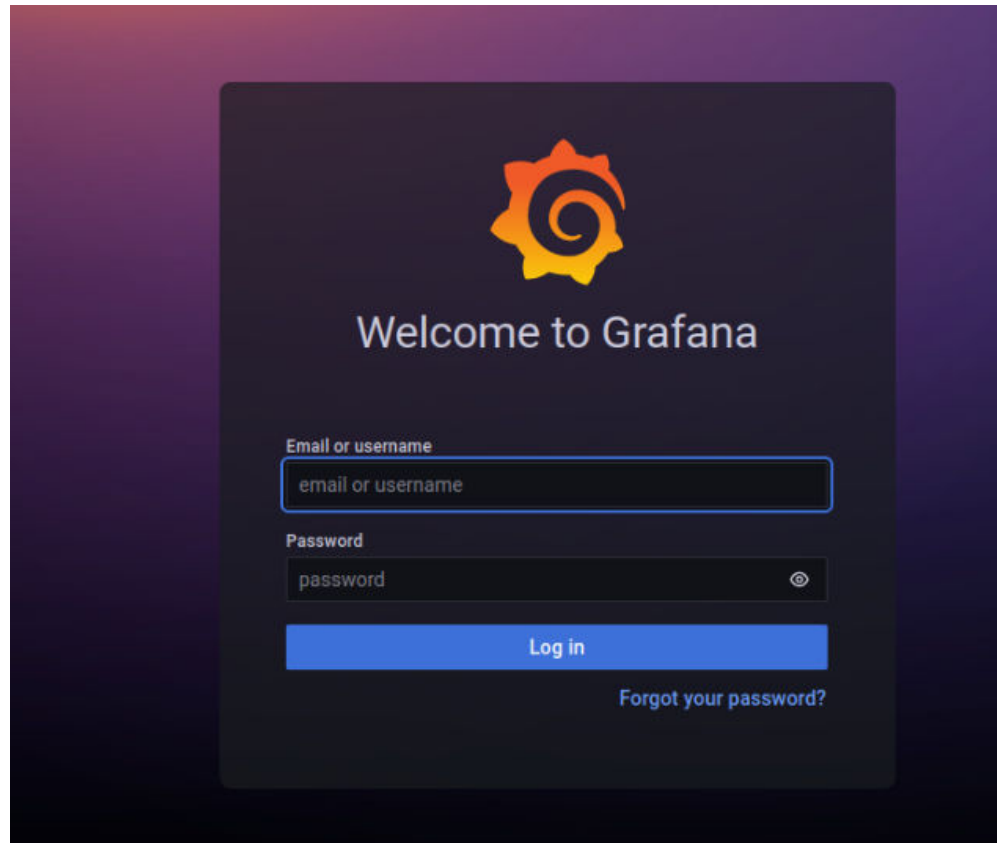
- a. Click the ECS name to go to the ECS details page. Then, click the **Security Groups** tab, and click **Manage Rule**.



- b. Click **Inbound Rules** and allow inbound traffic from TCP port 3000. By default, all outbound traffic is allowed.




- 7. Access **http://{EIP}:3000** in a browser. The default username and password for the first login are **admin**. After the login is successful, change the password as prompted.



### 4.2.2.3 Installing and Configuring Grafana on a Notebook Instance

#### Prerequisites

- A running CPU- or GPU-based notebook instance is available.
- A terminal is opened.



## Procedure

1. Run the following commands in sequence in your terminal to download and install Grafana:

```
mkdir -p /home/ma-user/work/grf
cd /home/ma-user/work/grf
wget https://dl.grafana.com/oss/release/grafana-9.1.6.linux-amd64.tar.gz
tar -zxvf grafana-9.1.6.linux-amd64.tar.gz
```

```
(PyTorch-1.8) [ma-user work]$ mkdir -p /home/ma-user/work/grf
(PyTorch-1.8) [ma-user work]$ cd /home/ma-user/work/grf
(PyTorch-1.8) [ma-user grf]$ wget https://dl.grafana.com/oss/release/grafana-9.1.6.linux-amd64.tar.gz
--2023-03-08 15:53:41-- https://dl.grafana.com/oss/release/grafana-9.1.6.linux-amd64.tar.gz
Resolving proxy.modelarts.com (proxy.modelarts.com)... 192.168.6.3
Connecting to proxy.modelarts.com (proxy.modelarts.com)|192.168.6.3|:80... connected.
Proxy request sent, awaiting response... 200 OK
Length: 81857482 (77M) [application/x-tar]
Saving to: 'grafana-9.1.6.linux-amd64.tar.gz.1'
grafana-9.1.6.linux-amd64.tar.gz.1 5K[====>] 4.41M 57.0KB/s eta 8m 19s
```

2. Register Grafana with jupyter-server-proxy.
  - a. Run the following commands in your terminal:

```
mkdir -p /home/ma-user/.local/etc/jupyter
vi /home/ma-user/.local/etc/jupyter/jupyter_notebook_config.py
```



```
(PyTorch-1.8) [ma-user grf]$mkdir -p /home/ma-user/.local/etc/jupyter
(PyTorch-1.8) [ma-user grf]$vi /home/ma-user/.local/etc/jupyter/jupyter_notebook_config.py
```

- b. In `jupyter_notebook_config.py`, add the following code, press **Esc** to exit, and type `:wq` to save the changes:

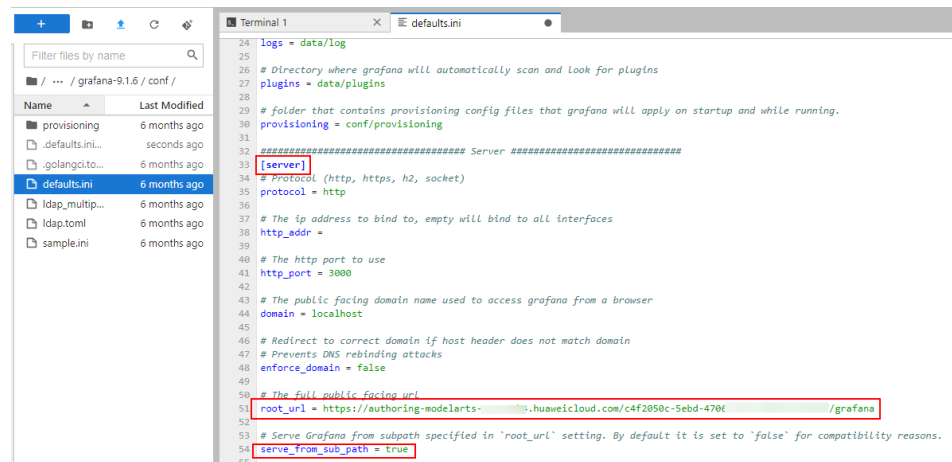
```
c.ServerProxy.servers = {
  'grafana': {
    'command': ['/home/ma-user/work/grf/grafana-9.1.6/bin/grafana-server', '--homepath', '/home/ma-user/work/grf/grafana-9.1.6', 'web'],
    'timeout': 1800,
    'port': 3000
  }
}
```

**NOTE**

If `jupyter_notebook_config.py` (path: `/home/ma-user/.local/etc/jupyter/jupyter_notebook_config.py`) contains the `c.ServerProxy.servers` field, add the corresponding key-value pair.


- 3. Modify the URL for accessing Grafana in JupyterLab.
  - a. In the navigation pane on the left, open the `vi /home/ma-user/work/grf/grafana-9.1.6/conf/defaults.ini` file.
  - b. Change the `root_url` and `serve_from_sub_path` fields in `[server]`.

**Figure 4-1** Modifying the `defaults.ini` file




In the file:

- The value of `root_url` is in the format of `https://{Jupyterlab domain name}/{Instance ID}/grafana`. You can obtain the domain name and instance ID from the address box of the JupyterLab page.



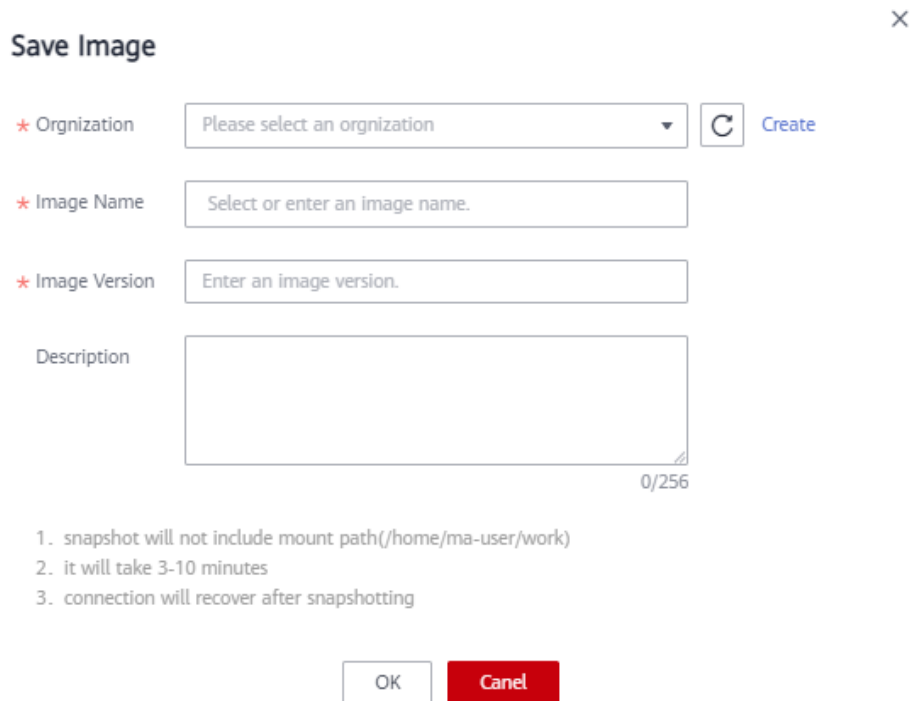
- Set `Serve_from_sub_path` to `true`.

- 4. Save the image of the notebook instance.
  - a. Log in to the ModelArts console and choose **DevEnviron > Notebook**. In the notebook instance list, choose **More > Save Image** in the **Operation** column of the target instance.




- b. In the **Save Image** dialog box, configure parameters. Click **OK** to save the image.

**Figure 4-2** Saving an image



- c. The image will be saved as a snapshot, and it will take about 5 minutes. During this period of time, do not perform any operations on the instance.

**Figure 4-3** Snapshotting



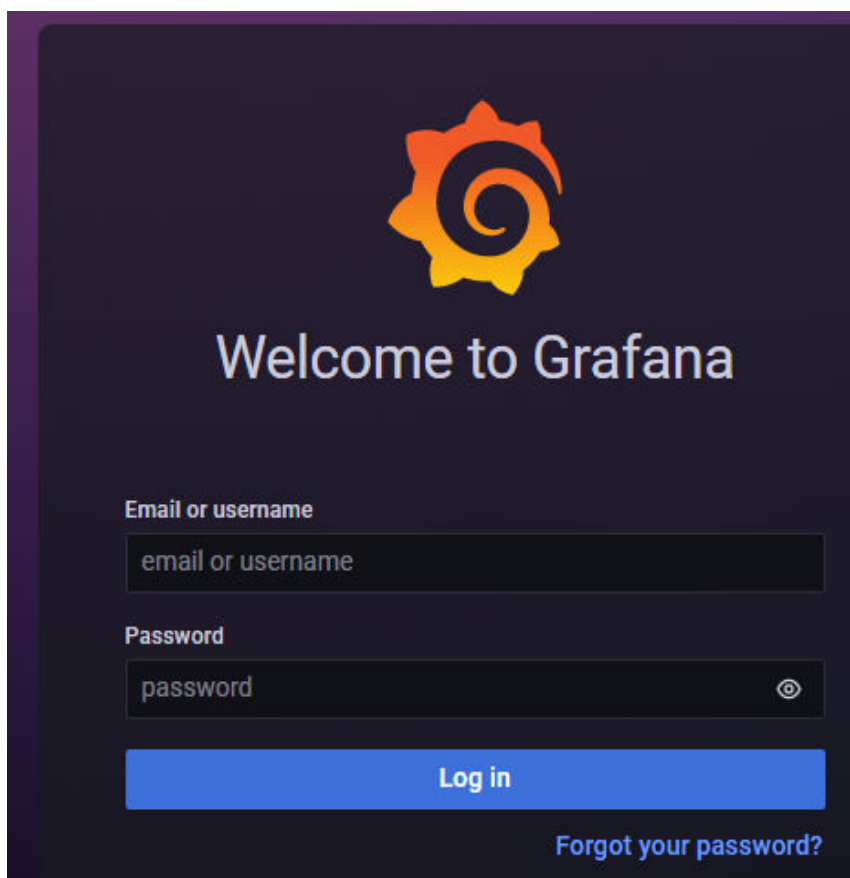
- d. After the image is saved, the instance status changes to **Running**. Then, restart the notebook instance.

Figure 4-4 Image saved

| Name            | Status              |
|-----------------|---------------------|
| notebc<br>295de | Running 59 minute.. |

5. Open the Grafana page.

Open a browser window and type the value of `root_url` configured in **3** in the address box. If the Grafana login page is displayed, Grafana is installed and configured in the notebook instance. The default username and password for the first login are **admin**. After the login is successful, change the password as prompted.



### 4.2.3 Configuring a Grafana Data Source

Before viewing ModelArts monitoring data on Grafana, configure the data source.


#### Prerequisites

- Grafana has been installed.

#### Procedure

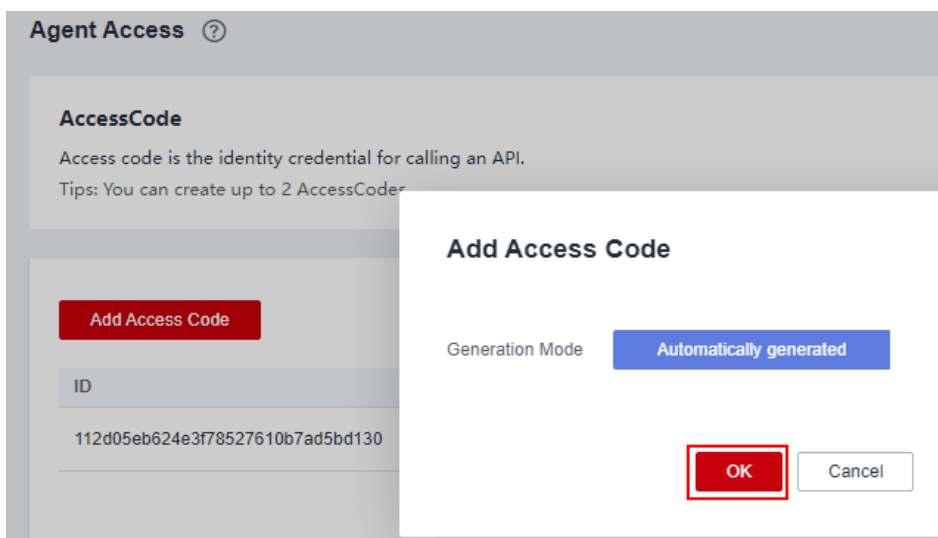
1. Add an access code.


- a. Log in to the AOM console.



- b. In the navigation pane on the left, choose **Configuration Management > Agent Access**, and click **Add Access Code** to generate an access code.

**Figure 4-5** Generating an access code



- c. Click  to view the generated access code.

**Figure 4-6** Viewing the access code

| ID       | AccessCode   |
|----------|--|
| 2a406bl  |  X****b |
| 624ce2ff |  e****G |

2. Obtain the data source URL.  
The URL is in the format of **https://{Endpoint}/v1/{project\_id}**.
  - You can obtain the AOM endpoint information from **Regions and Endpoints**.
  - Set **project\_id** to the project ID of the corresponding region. You can obtain the project ID from **My Credentials**.

Figure 4-7 My Credentials

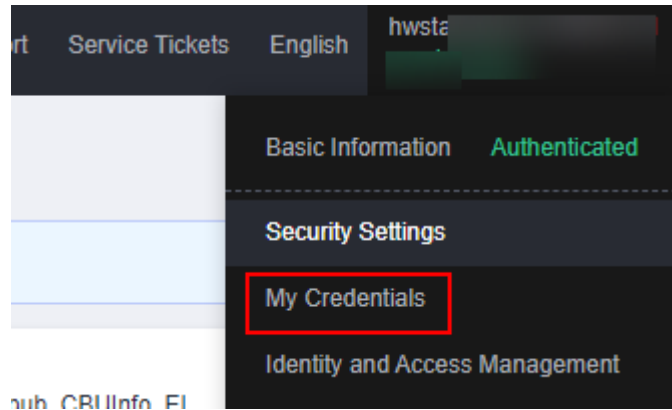
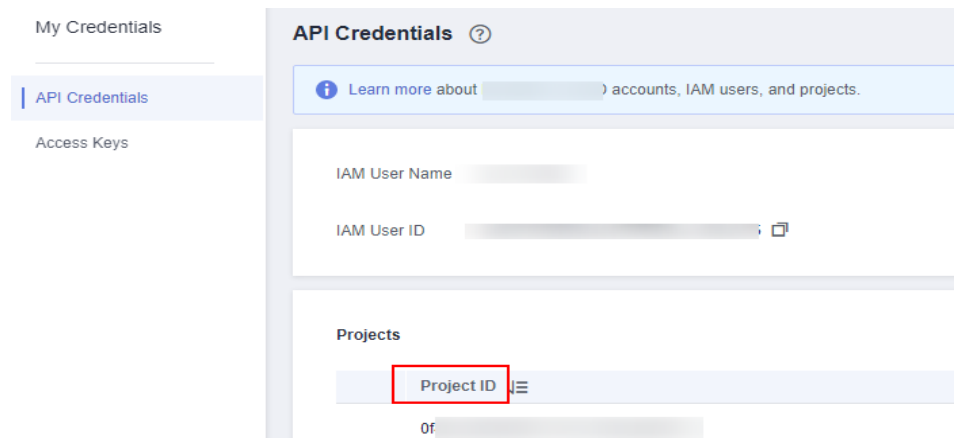
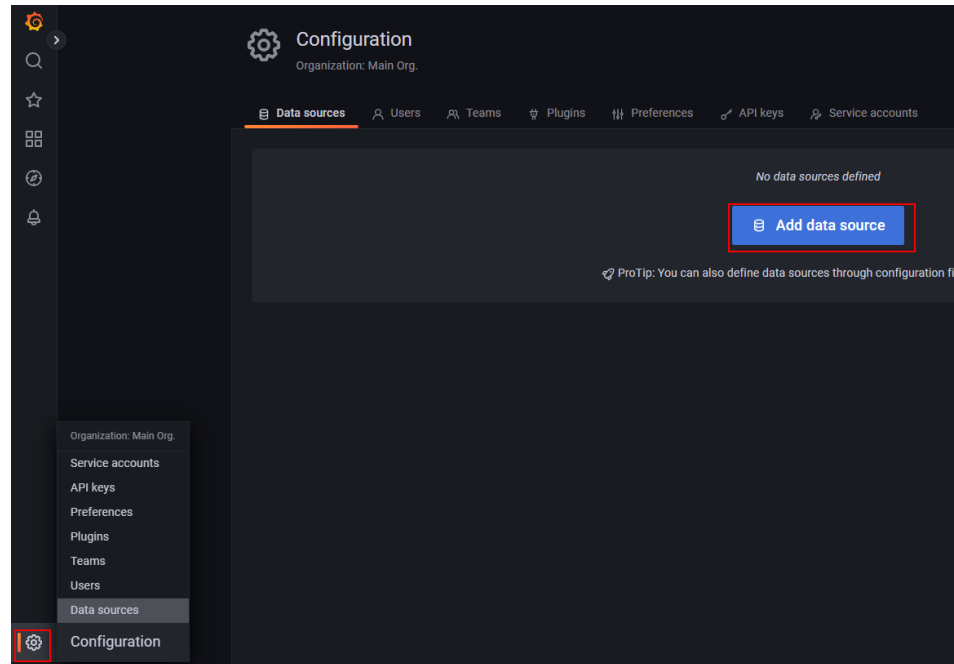


Figure 4-8 Obtaining the project ID



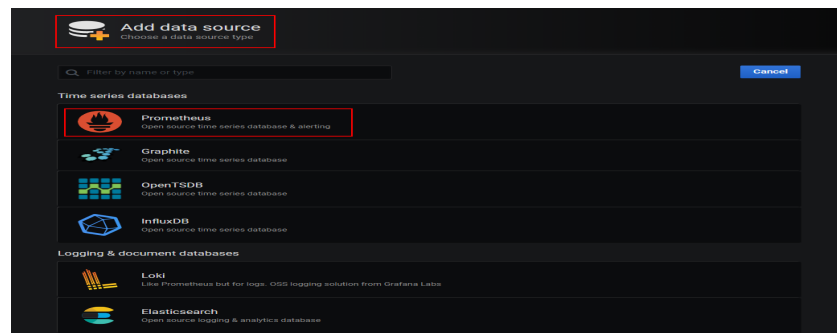
3. Add a data source to Grafana.
  - a. Log in to Grafana. The default username and password for the first login are **admin**. After the login is successful, change the password as prompted.
  - b. In the navigation pane, choose **Configuration > Data Sources**. Then, click **Add data source**.

Figure 4-9 Configuring Grafana



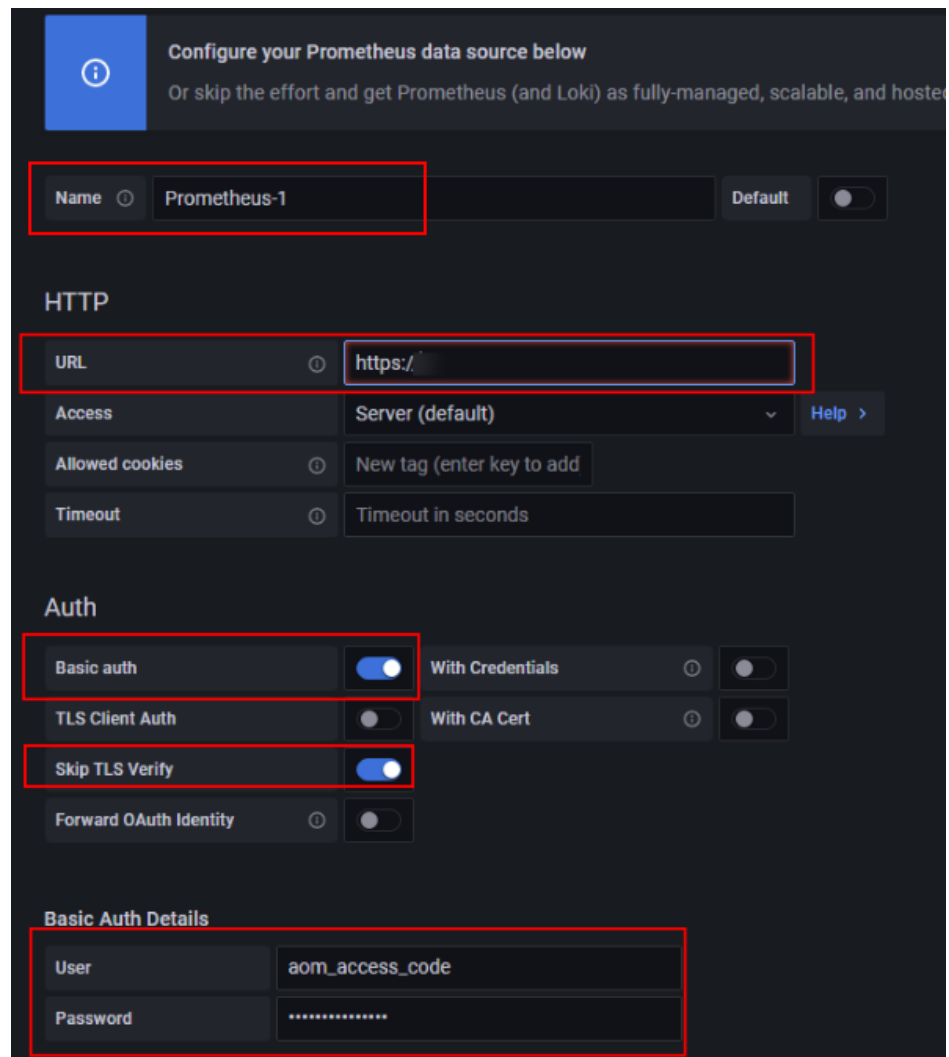
- c. Click **Prometheus** to access the configuration page.

Figure 4-10 Entering the Prometheus configuration page



- d. Configure parameters as shown in the following figure.

**Figure 4-11** Configuring a Grafana data source



**NOTE**

The actual Grafana version varies depending on the installation method. [Figure 4-11](#) is only an example.

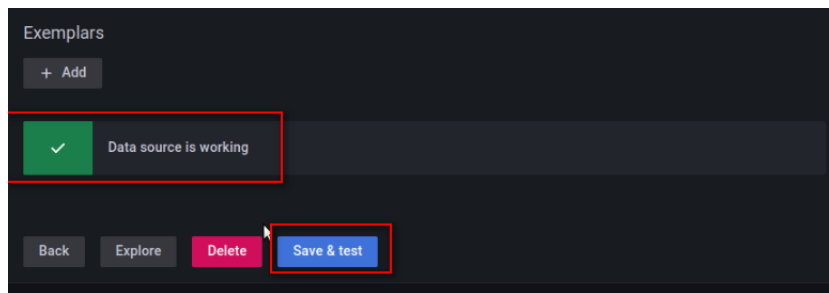
**Table 4-1** Parameters

| Parameter       | Description  |
|-----------------|--|
| Name            | Customizable name  |
| URL             | URL <code>https://{Endpoint}/v1/{project_id}</code> combined in <a href="#">Obtain the data source URL</a> . |
| Basic auth      | Enabled  |
| Skip TLS Verify | Enabled  |
| User            | aom_access_code  |

| Parameter | Description   |
|-----------|---|
| Password  | Access code generated in <a href="#">Add an access code</a> . |

- e. After the configuration, click **Save & test**. If the message **Data source is working** is displayed, the data source is configured.

**Figure 4-12** Data source added



## 4.2.4 Using Grafana to Configure Dashboards and View Metric Data

In Grafana, you can customize dashboards for various views. ModelArts also provides configuration templates for clusters. This section describes how to configure a dashboard by using a ModelArts template or creating a dashboard. For more usage, see [Grafana tutorials](#).

### Preparations

ModelArts provides templates for cluster view, node view, user view, task view, and task details view. These templates can be downloaded from Grafana official documents. You can import and use them on Dashboards.


**Table 4-2** Template download URLs

| Template Name     | Download URL  |
|-------------------|---|
| Cluster view      | <a href="https://grafana.com/grafana/dashboards/18582-modelarts-cluster-view/">https://grafana.com/grafana/dashboards/18582-modelarts-cluster-view/</a>         |
| Node view         | <a href="https://grafana.com/grafana/dashboards/18583-modelarts-node-view/">https://grafana.com/grafana/dashboards/18583-modelarts-node-view/</a>               |
| User view         | <a href="https://grafana.com/grafana/dashboards/18588-modelarts-user-view/">https://grafana.com/grafana/dashboards/18588-modelarts-user-view/</a>               |
| Task view         | <a href="https://grafana.com/grafana/dashboards/18604-modelarts-task-view/">https://grafana.com/grafana/dashboards/18604-modelarts-task-view/</a>               |
| Task details view | <a href="https://grafana.com/grafana/dashboards/18590-modelarts-task-detail-view/">https://grafana.com/grafana/dashboards/18590-modelarts-task-detail-view/</a> |



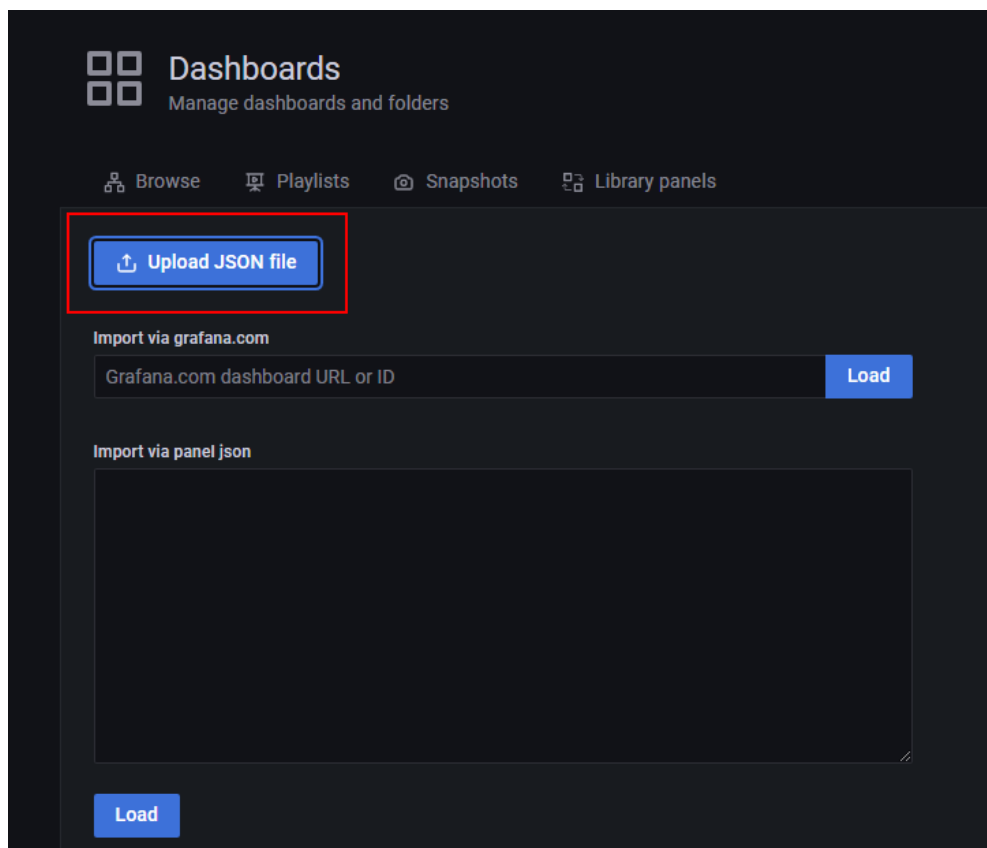
## Using a ModelArts Template to View Metrics

1. (Optional) Select the template you want to use. **Preparations** displays the download addresses of all templates. Open the target address and click **Download JSON**.
2. Open **Dashboards** and choose **New > Import**.

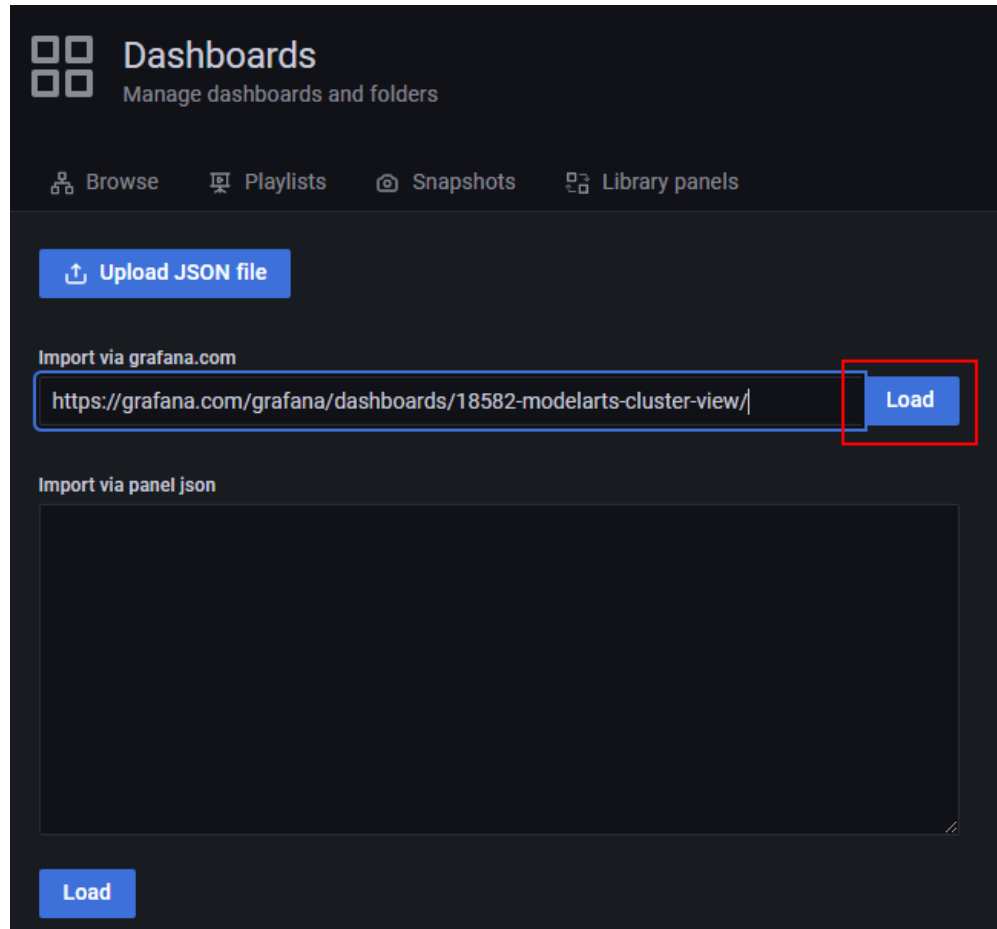


3. Import a dashboard template in either of the following ways:
  - Method 1: Upload the JSON file downloaded in 1, as shown in **Figure 4-13**.
  - Method 2: Copy the template download address provided in **Preparations** and click **Load**, as shown in **Figure 4-14**.

**Figure 4-13** Uploading a JSON file to import a dashboard template

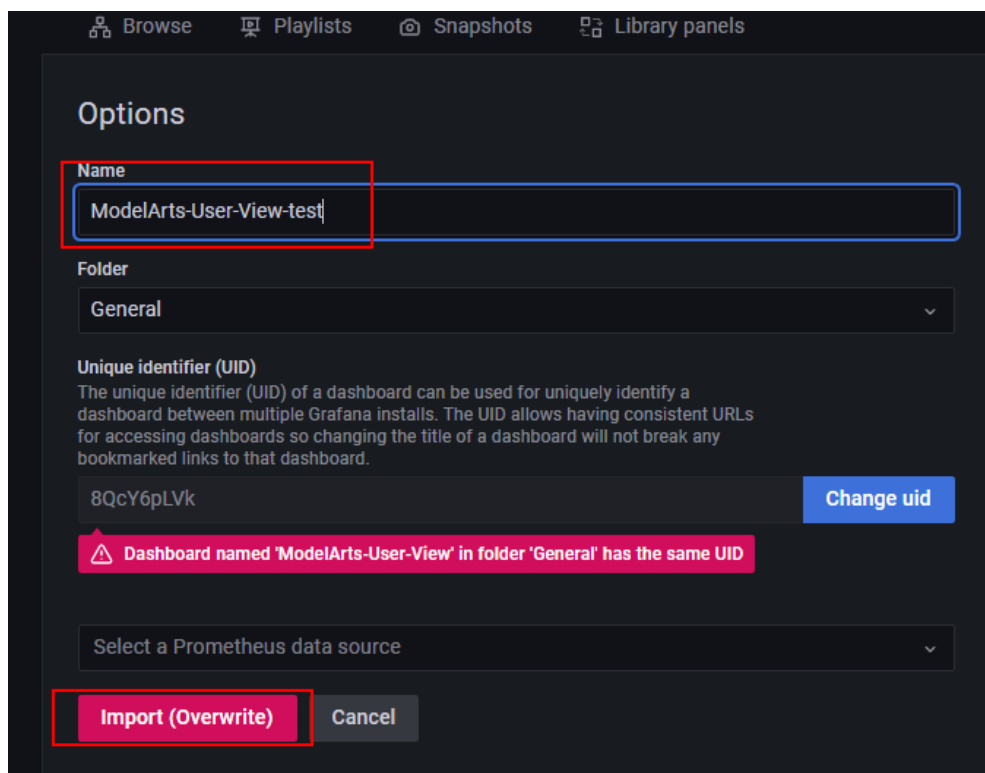


**Figure 4-14** Copying the template address and importing the dashboard template



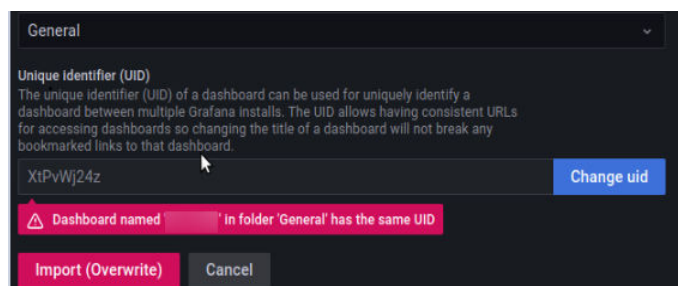
4. Change the view name and click **Import**.

Figure 4-15 Changing the view name

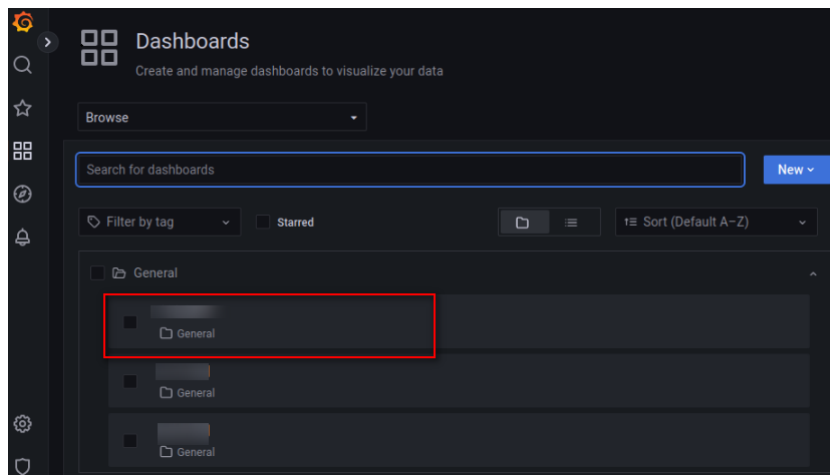


Note: If a message is displayed, indicating that the UID is duplicate, change the UID in the JSON file and click **Import**.

Figure 4-16 Changing the UID



5. After the import, view the imported views in **Dashboards**. Then, click a view to open the monitoring page.

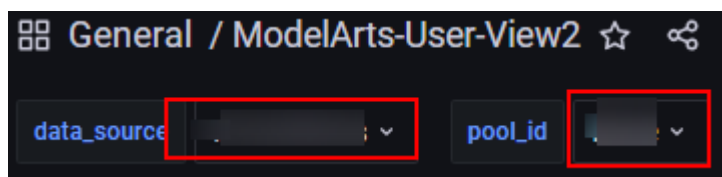


6. Use the template.

After the import is successful, you can click the template to view its details. This section introduces some common functions.

- Changing the data source and resource pool

**Figure 4-17** Changing the data source and resource pool



Click the area marked by the red box. A drop-down list will appear. From there, you can change the data source and the resource pool.

- Refreshing data



Click the refresh button in the upper right corner to refresh all data on the dashboard. The data on each panel is also updated.

- Changing the automatic refresh time

The default refresh interval of a template is 15 minutes. If you need to update the interval, change the value from the drop-down list box in the upper right corner.

- Changing the time range for obtaining dashboard data

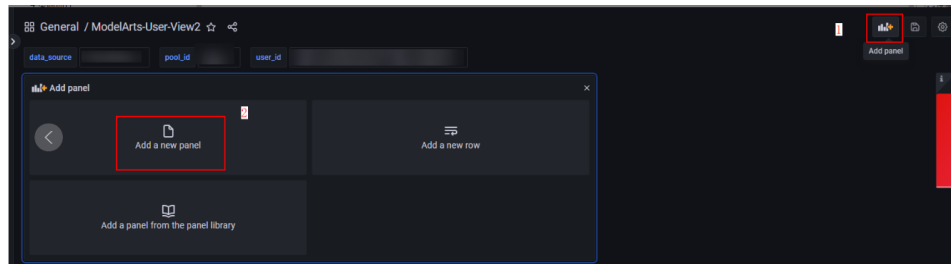
**Figure 4-18** Changing the time range for obtaining data



Click the button in the upper right corner to change time range for obtaining data. This time range affects all panels except those with a fixed time.

- Adding a panel

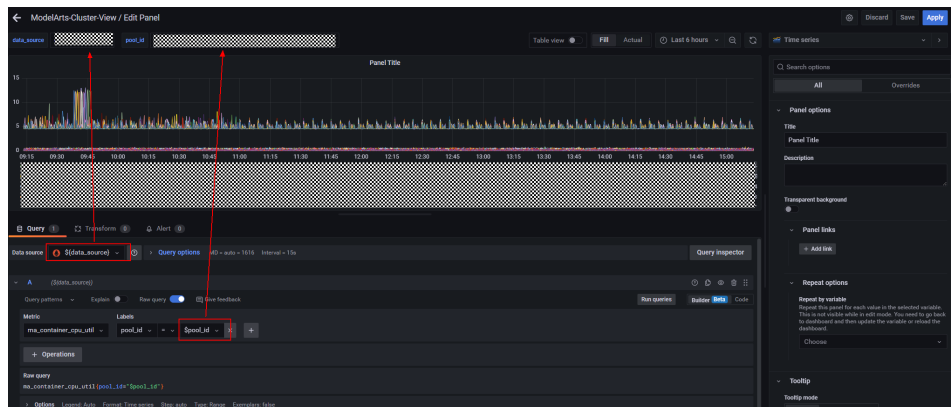
Figure 4-19 Adding a panel



Click the + icon in the upper right corner to add a panel.

After a panel is added, you can obtain the data in the panel. Configure the data source and resource pool as follows to use the current dashboard settings.

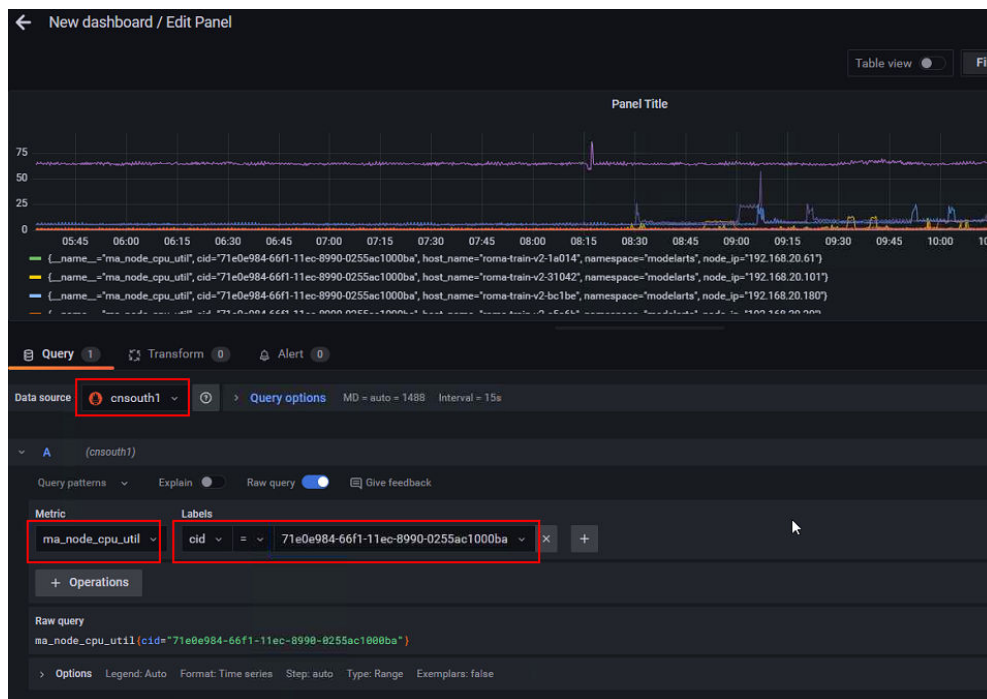
Figure 4-20 Using the current dashboard settings



## Creating a Dashboard to View Metrics

1. Open **Dashboards**, click **New**, and choose **New Dashboard**.
2. Click **Add a new panel**.
3. On the **New dashboard / Edit Panel** page, set the following parameters:  
**Data source:** [Configured Grafana data source](#)  
**Metric:** Metric name. You can obtain the metric to be queried by referring to [Table 4-3](#), [Table 4-4](#), and [Table 4-5](#).  
**Labels:** Used for filtering the metric. For details, see [Table 4-6](#).

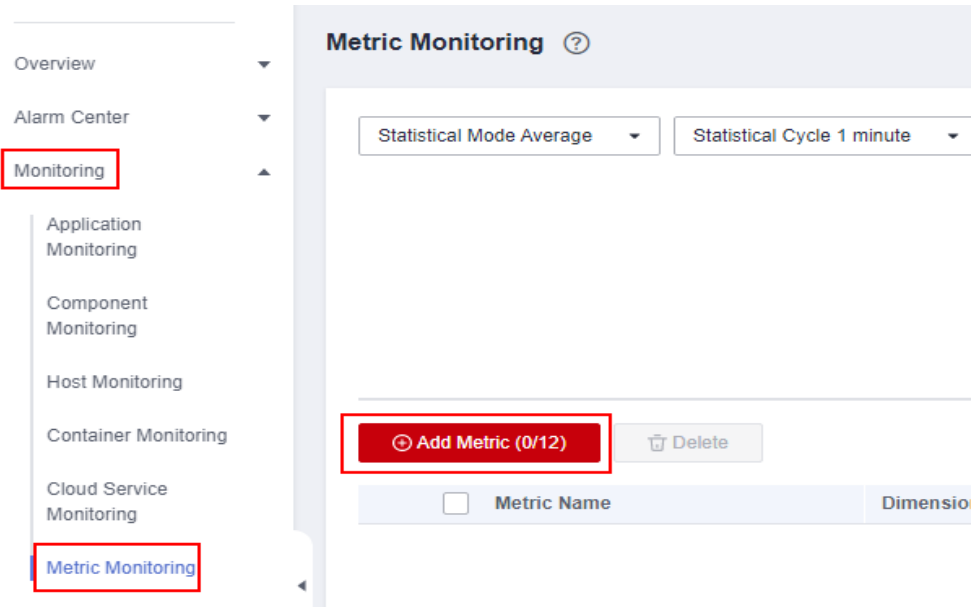
Figure 4-21 Creating a dashboard to view metrics



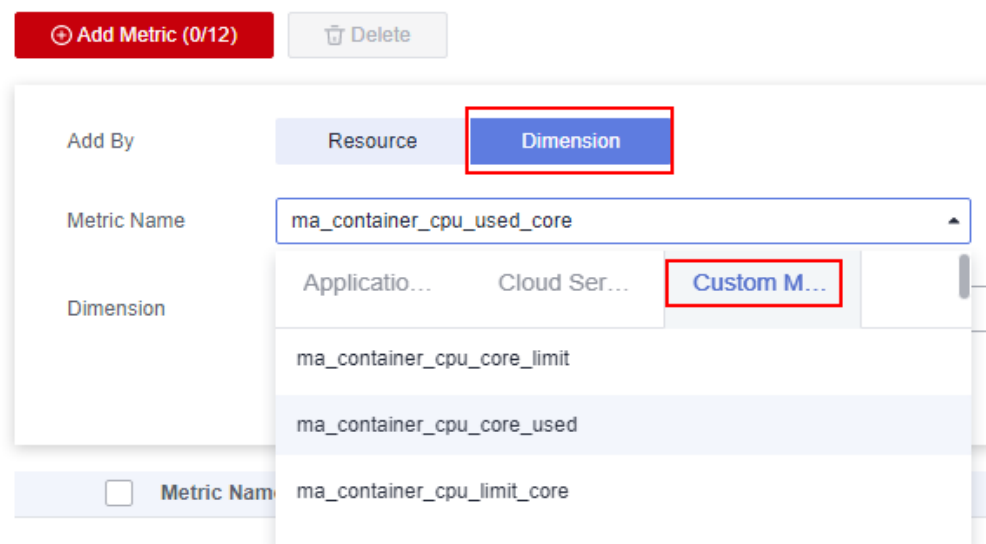
## 4.3 Viewing All ModelArts Monitoring Metrics on the AOM Console

ModelArts periodically collects the usage of key metrics (such as GPUs, NPUs, CPUs, and memory) of each node in a resource pool as well as the usage of key metrics of the development environment, training jobs, and inference services, and reports the data to AOM. You can view the information on AOM.

1. Log in to the console and search for **AOM** to go to the AOM console.
2. Choose **Monitoring > Metric Monitoring**. On the **Metric Monitoring** page that is displayed, click **Add Metric**.



3. Add metrics and click Confirm.



- **Add By:** Select Dimension.
- **Metric Name:** Click **Custom Metrics**. Select the desired ones for query. For details, see [Table 4-3](#), [Table 4-4](#), and [Table 4-5](#).
- **Dimension:** Enter the tag for filtering the metric. For details, see [Table 4-6](#). The following shows an example.

Add By:

Metric Name:

Dimension:

Metric Name      Dimensions

4. View the metrics.



Table 4-3 Container metrics

| Category | Name                  | Metric                                 | Description  | Unit  | Value Range |
|----------|-----------------------|--|--|-------|-------------|
| CPU      | CPU Usage             | ma_container_cpu_util                  | CPU usage of a measured object   | %     | 0%–100%     |
|          | Used CPU Cores        | ma_container_cpu_used_core             | Number of CPU cores used by a measured object                          | Cores | ≥ 0         |
|          | Total CPU Cores       | ma_container_cpu_limit_core            | Total number of CPU cores that have been applied for a measured object | Cores | ≥ 1         |
| Memory   | Total Physical Memory | ma_container_memory_capacity_megabytes | Total physical memory that has been applied for a measured object      | MB    | ≥ 0         |



| Category   | Name                  | Metric                                  | Description  | Unit | Value Range |
|------------|-----------------------|---|--|------|-------------|
|            | Physical Memory Usage | ma_container_memory_util                | Percentage of the used physical memory to the total physical memory  | %    | 0%–100%     |
|            | Used Physical Memory  | ma_container_memory_used_megabytes      | Physical memory that has been used by a measured object ( <b>container_memory_working_set_bytes</b> in the current working set)<br><br>(Memory usage in a working set = Active anonymous page and cache, and file-baked page ≤ <b>container_memory_usage_bytes</b> ) | MB   | ≥ 0         |
| Storage    | Disk Read Rate        | ma_container_disk_read_kilobytes        | Volume of data read from a disk per second   | KB/s | ≥ 0         |
|            | Disk Write Rate       | ma_container_disk_write_kilobytes       | Volume of data written into a disk per second  | KB/s | ≥ 0         |
| GPU memory | Total GPU Memory      | ma_container_gpu_memory_total_megabytes | Total GPU memory of a training job   | MB   | > 0         |
|            | GPU Memory Usage      | ma_container_gpu_memory_util            | Percentage of the used GPU memory to the total GPU memory  | %    | 0%–100%     |
|            | Used GPU Memory       | ma_container_gpu_memory_used_megabytes  | GPU memory used by a measured object   | MB   | ≥ 0         |

| Category    | Name                       | Metric                               | Description   | Unit       | Value Range    |
|-------------|----------------------------|--------------------------------------|---|------------|----------------|
| GPU         | GPU Usage                  | ma_container_gpu_util                | GPU usage of a measured object  | %          | 0%–100%        |
|             | GPU Memory Bandwidth Usage | ma_container_gpu_mem_copy_util       | GPU memory bandwidth usage of a measured object For example, the maximum memory bandwidth of NVIDIA GPU V100 is 900 GB/s. If the current memory bandwidth is 450 GB/s, the memory bandwidth usage is 50%. | %          | 0%–100%        |
|             | GPU Encoder Usage          | ma_container_gpu_encode_util         | GPU encoder usage of a measured object  | %          | %              |
|             | GPU Decoder Usage          | ma_container_gpu_decode_util         | GPU decoder usage of a measured object  | %          | %              |
|             | GPU Temperature            | DCGM_FI_DEV_GPU_TEMP                 | GPU temperature   | °C         | Natural number |
|             | GPU Power                  | DCGM_FI_DEV_POWER_USAGE              | GPU power   | Watt (W)   | > 0            |
|             | GPU Memory Temperature     | DCGM_FI_DEV_MEMORY_TEMP              | GPU memory temperature  | °C         | Natural number |
| Network I/O | Downlink Rate (BPS)        | ma_container_network_receive_bytes   | Inbound traffic rate of a measured object   | Bytes/s    | ≥ 0            |
|             | Downlink Rate (PPS)        | ma_container_network_receive_packets | Number of data packets received by a NIC per second   | Packet s/s | ≥ 0            |

| Category                 | Name                                 | Metric                                      | Description   | Unit       | Value Range |
|--------------------------|--------------------------------------|---|---|------------|-------------|
|                          | Downlink Error Rate                  | ma_container_network_receive_error_packets  | Number of error packets received by a NIC per second  | Packet s/s | ≥ 0         |
|                          | Uplink Rate (BPS)                    | ma_container_network_transmit_bytes         | Outbound traffic rate of a measured object  | Bytes/s    | ≥ 0         |
|                          | Uplink Error Rate                    | ma_container_network_transmit_error_packets | Number of error packets sent by a NIC per second  | Packet s/s | ≥ 0         |
|                          | Uplink Rate (PPS)                    | ma_container_network_transmit_packets       | Number of data packets sent by a NIC per second   | Packet s/s | ≥ 0         |
| Notebook service metrics | Notebook Cache Directory Size        | ma_container_notebook_cache_dir_size_bytes  | A high-speed local disk is attached to the / <b>cache</b> directory for GPU notebook instances. This metric indicates the total size of the directory.  | Bytes      | ≥ 0         |
|                          | Notebook Cache Directory Utilization | ma_container_notebook_cache_dir_util        | A high-speed local disk is attached to the / <b>cache</b> directory for GPU notebook instances. This metric indicates the utilization of the directory. | %          | 0%–100%     |

**Table 4-4** Node metrics (collected only in dedicated resource pools)

| Category | Name                  | Metric                          | Description  | Unit    | Value Range |
|----------|-----------------------|---------------------------------|--|---------|-------------|
| CPU      | Total CPU Cores       | ma_node_cpu_limit_core          | Total number of CPU cores that have been applied for a measured object | Cores   | ≥ 1         |
|          | Used CPU Cores        | ma_node_cpu_used_core           | Number of CPU cores used by a measured object                          | Cores   | ≥ 0         |
|          | CPU Usage             | ma_node_cpu_util                | CPU usage of a measured object   | %       | 0%–100%     |
|          | CPU I/O Wait Time     | ma_node_cpu_iowait_counter      | Disk I/O wait time accumulated since system startup                    | jiffies | ≥ 0         |
| Memory   | Physical Memory Usage | ma_node_memory_util             | Percentage of the used physical memory to the total physical memory    | %       | 0%–100%     |
|          | Total Physical Memory | ma_node_memory_total_mega bytes | Total physical memory that has been applied for a measured object      | MB      | ≥ 0         |

| Category    | Name                | Metric                                      | Description   | Unit    | Value Range |
|-------------|---------------------|---|---|---------|-------------|
| Network I/O | Downlink Rate (BPS) | ma_node_network_receive_rate_bytes_seconds  | Inbound traffic rate of a measured object   | Bytes/s | ≥ 0         |
|             | Uplink Rate (BPS)   | ma_node_network_transmit_rate_bytes_seconds | Outbound traffic rate of a measured object  | Bytes/s | ≥ 0         |
| Storage     | Disk Read Rate      | ma_node_disk_read_rate_kilobytes_seconds    | Volume of data read from a disk per second (Only data disks used by containers are collected.)    | KB/s    | ≥ 0         |
|             | Disk Write Rate     | ma_node_disk_write_rate_kilobytes_seconds   | Volume of data written into a disk per second (Only data disks used by containers are collected.) | KB/s    | ≥ 0         |
|             | Total Cache         | ma_node_cache_space_capacity_megabytes      | Total cache of the Kubernetes space   | MB      | ≥ 0         |
|             | Used Cache          | ma_node_cache_space_used_capacity_megabytes | Used cache of the Kubernetes space  | MB      | ≥ 0         |

| Category | Name                  | Metric  | Description                                      | Unit    | Value Range |
|----------|-----------------------|---|--|---------|-------------|
|          | Total Container Space | ma_node_container_space_capacity_megabytes      | Total container space                            | MB      | ≥ 0         |
|          | Used Container Space  | ma_node_container_space_used_capacity_megabytes | Used container space                             | MB      | ≥ 0         |
|          | Disk Information      | ma_node_disk_info                               | Basic disk information                           | N/A     | ≥ 0         |
|          | Total Reads           | ma_node_disk_reads_completed_total              | Total number of successful reads                 | N/A     | ≥ 0         |
|          | Merged Reads          | ma_node_disk_reads_merged_total                 | Number of merged reads                           | N/A     | ≥ 0         |
|          | Bytes Read            | ma_node_disk_read_bytes_total                   | Total number of bytes that are successfully read | Bytes   | ≥ 0         |
|          | Read Time Spent       | ma_node_disk_read_time_seconds_total            | Time spent on all reads                          | Seconds | ≥ 0         |
|          | Total Writes          | ma_node_disk_writes_completed_total             | Total number of successful writes                | N/A     | ≥ 0         |
|          | Merged Writes         | ma_node_disk_writes_merged_total                | Number of merged writes                          | N/A     | ≥ 0         |

| Category | Name                        | Metric                                      | Description   | Unit    | Value Range |
|----------|-----------------------------|---|---|---------|-------------|
|          | Written Bytes               | ma_node_disk_written_bytes_total            | Total number of bytes that are successfully written       | Bytes   | ≥ 0         |
|          | Write Time Spent            | ma_node_disk_write_time_seconds_total       | Time spent on all write operations                        | Seconds | ≥ 0         |
|          | Ongoing I/Os                | ma_node_disk_io_now                         | Number of ongoing I/Os                                    | N/A     | ≥ 0         |
|          | I/O Execution Duration      | ma_node_disk_io_time_seconds_total          | Time spent on executing I/Os                              | Seconds | ≥ 0         |
|          | I/O Execution Weighted Time | ma_node_disk_io_time_weighted_seconds_total | The weighted number of seconds spent doing I/Os           | Seconds | ≥ 0         |
| GPU      | GPU Usage                   | ma_node_gpu_util                            | GPU usage of a measured object                            | %       | 0%–100%     |
|          | Total GPU Memory            | ma_node_gpu_memory_total_megabytes          | Total GPU memory of a measured object                     | MB      | > 0         |
|          | GPU Memory Usage            | ma_node_gpu_memory_util                     | Percentage of the used GPU memory to the total GPU memory | %       | 0%–100%     |

| Category                   | Name                                   | Metric  | Description  | Unit                                | Value Range    |
|----------------------------|--|---|--|-------------------------------------|----------------|
|                            | Used GPU Memory                        | ma_node_gpu_mem_used_megabytes                    | GPU memory used by a measured object   | MB                                  | $\geq 0$       |
|                            | Tasks on a Shared GPU                  | node_gpu_share_job_count                          | Number of tasks running on a shared GPU  | Number                              | $\geq 0$       |
|                            | GPU Temperature                        | DCGM_FI_DEV_GPU_TEMP                              | GPU temperature  | °C                                  | Natural number |
|                            | GPU Power                              | DCGM_FI_DEV_POWER_USAGE                           | GPU power  | Watt (W)                            | $> 0$          |
|                            | GPU Memory Temperature                 | DCGM_FI_DEV_MEMORY_TEMP                           | GPU memory temperature   | °C                                  | Natural number |
| InfiniBand or RoCE network | Total Amount of Data Received by a NIC | ma_node_infiniband_port_received_data_bytes_total | The total number of data octets, divided by 4, (counting in double words, 32 bits), received on all VLS from the port. | (counting in double words, 32 bits) | $\geq 0$       |



| Category | Name                               | Metric   | Description   | Unit                                | Value Range |
|----------|------------------------------------|--|---|-------------------------------------|-------------|
|          | Total Amount of Data Sent by a NIC | ma_node_infiniband_port_transmitted_data_bytes_total | The total number of data octets, divided by 4, (counting in double words, 32 bits), transmitted on all VLs from the port. | (counting in double words, 32 bits) | $\geq 0$    |

| Category            | Name                        | Metric                                  | Description   | Unit | Value Range |
|---------------------|-----------------------------|---|---|------|-------------|
| NFS mounting status | NFS Getattr Congestion Time | ma_node_mountstats_getattr_backlog_wait | Getattr is an NFS operation that retrieves the attributes of a file or directory, such as size, permissions, owner, etc. Backlog wait is the time that the NFS requests have to wait in the backlog queue before being sent to the NFS server. It indicates the congestion on the NFS client side. A high backlog wait can cause poor NFS performance and slow system response times. | ms   | ≥ 0         |

| Category | Name                                 | Metric                                     | Description  | Unit | Value Range |
|----------|--------------------------------------|--|--|------|-------------|
|          | NFS<br>Getattr<br>Round Trip<br>Time | ma_node_<br>mountstat<br>s_getattr_r<br>tt | <p>Getattr is an NFS operation that retrieves the attributes of a file or directory, such as size, permissions, owner, etc.</p> <p>RTT stands for Round Trip Time and it is the time from when the kernel RPC client sends the RPC request to the time it receives the reply<sup>34</sup>. RTT includes network transit time and server execution time. RTT is a good measurement for NFS latency. A high RTT can indicate network or server issues.</p> | ms   | ≥ 0         |

| Category | Name                       | Metric                                 | Description   | Unit | Value Range |
|----------|----------------------------|--|---|------|-------------|
|          | NFS Access Congestion Time | ma_node_mountstats_access_backlog_wait | Access is an NFS operation that checks the access permissions of a file or directory for a given user. Backlog wait is the time that the NFS requests have to wait in the backlog queue before being sent to the NFS server. It indicates the congestion on the NFS client side. A high backlog wait can cause poor NFS performance and slow system response times. | ms   | ≥ 0         |

| Category | Name                       | Metric                        | Description  | Unit | Value Range |
|----------|----------------------------|-------------------------------|--|------|-------------|
|          | NFS Access Round Trip Time | ma_node_mountstats_access_rtt | Access is an NFS operation that checks the access permissions of a file or directory for a given user. RTT stands for Round Trip Time and it is the time from when the kernel RPC client sends the RPC request to the time it receives the reply <sup>34</sup> . RTT includes network transit time and server execution time. RTT is a good measurement for NFS latency. A high RTT can indicate network or server issues. | ms   | ≥ 0         |

| Category | Name                       | Metric                                 | Description  | Unit | Value Range |
|----------|----------------------------|--|--|------|-------------|
|          | NFS Lookup Congestion Time | ma_node_mountstats_lookup_backlog_wait | Lookup is an NFS operation that resolves a file name in a directory to a file handle. Backlog wait is the time that the NFS requests have to wait in the backlog queue before being sent to the NFS server. It indicates the congestion on the NFS client side. A high backlog wait can cause poor NFS performance and slow system response times. | ms   | ≥ 0         |

| Category | Name                       | Metric                        | Description   | Unit | Value Range |
|----------|----------------------------|-------------------------------|---|------|-------------|
|          | NFS Lookup Round Trip Time | ma_node_mountstats_lookup_rtt | Lookup is an NFS operation that resolves a file name in a directory to a file handle. RTT stands for Round Trip Time and it is the time from when the kernel RPC client sends the RPC request to the time it receives the reply <sup>34</sup> . RTT includes network transit time and server execution time. RTT is a good measurement for NFS latency. A high RTT can indicate network or server issues. | ms   | ≥ 0         |

| Category | Name                     | Metric                               | Description  | Unit | Value Range |
|----------|--------------------------|--------------------------------------|--|------|-------------|
|          | NFS Read Congestion Time | ma_node_mountstats_read_backlog_wait | Read is an NFS operation that reads data from a file. Backlog wait is the time that the NFS requests have to wait in the backlog queue before being sent to the NFS server. It indicates the congestion on the NFS client side. A high backlog wait can cause poor NFS performance and slow system response times. | ms   | ≥ 0         |



| Category | Name                     | Metric                      | Description   | Unit | Value Range |
|----------|--------------------------|-----------------------------|---|------|-------------|
|          | NFS Read Round Trip Time | ma_node_mountstats_read_rtt | Read is an NFS operation that reads data from a file. RTT stands for Round Trip Time and it is the time from when the kernel RPC client sends the RPC request to the time it receives the reply <sup>34</sup> . RTT includes network transit time and server execution time. RTT is a good measurement for NFS latency. A high RTT can indicate network or server issues. | ms   | ≥ 0         |

| Category | Name                      | Metric                                | Description  | Unit | Value Range |
|----------|---------------------------|---------------------------------------|--|------|-------------|
|          | NFS Write Congestion Time | ma_node_mountstats_write_backlog_wait | Write is an NFS operation that writes data to a file. Backlog wait is the time that the NFS requests have to wait in the backlog queue before being sent to the NFS server. It indicates the congestion on the NFS client side. A high backlog wait can cause poor NFS performance and slow system response times. | ms   | ≥ 0         |

| Category | Name                      | Metric                       | Description   | Unit | Value Range |
|----------|---------------------------|------------------------------|---|------|-------------|
|          | NFS Write Round Trip Time | ma_node_mountstats_write_rtt | Write is an NFS operation that writes data to a file. RTT stands for Round Trip Time and it is the time from when the kernel RPC client sends the RPC request to the time it receives the reply <sup>34</sup> . RTT includes network transit time and server execution time. RTT is a good measurement for NFS latency. A high RTT can indicate network or server issues. | ms   | ≥ 0         |

**Table 4-5** Diagnosis (InfiniBand, collected only in dedicated resource pools)

| Category                   | Name                     | Metric                                       | Description   | Unit        | Value Range    |
|----------------------------|--------------------------|--|---|-------------|----------------|
| InfiniBand or RoCE network | PortXmitData             | infiniband_port_xmit_data_total              | The total number of data octets, divided by 4, (counting in double words, 32 bits), transmitted on all VLs from the port. | Total count | Natural number |
|                            | PortRcvData              | infiniband_port_rcv_data_total               | The total number of data octets, divided by 4, (counting in double words, 32 bits), received on all VLs from the port.    | Total count | Natural number |
|                            | SymbolErrorCounter       | infiniband_symbol_error_counter_total        | Total number of minor link errors detected on one or more physical lanes.   | Total count | Natural number |
|                            | LinkErrorRecoveryCounter | infiniband_link_error_recovery_counter_total | Total number of times the Port Training state machine has successfully completed the link error recovery process.         | Total count | Natural number |

| Category | Name                        | Metric   | Description   | Unit        | Value Range    |
|----------|-----------------------------|--|---|-------------|----------------|
|          | PortRcvErrors               | infiniband_port_rcv_errors_total                 | Total number of packets containing errors that were received on the port including:<br><br>Local physical errors (ICRC, VCRC, LPCRC, and all physical errors that cause entry into the BAD PACKET or BAD PACKET DISCARD states of the packet receiver state machine)<br><br>Malformed data packet errors (LVer, length, VL)<br><br>Malformed link packet errors (operand, length, VL)<br><br>Packets discarded due to buffer overrun (overflow) | Total count | Natural number |
|          | LocalLinkIntegrityErrors    | infiniband_local_link_integrity_errors_total     | This counter indicates the number of retries initiated by a link transfer layer receiver.   | Total count | Natural number |
|          | PortRcvRemotePhysicalErrors | infiniband_port_rcv_remote_physical_errors_total | Total number of packets marked with the EBP delimiter received on the port.   | Total count | Natural number |
|          | PortRcvSwitchRelayErrors    | infiniband_port_rcv_switch_relay_errors_total    | Total number of packets received on the port that were discarded when they could not be forwarded by the switch relay for the following reasons:<br><br>DLID mapping<br><br>VL mapping<br><br>Looping (output port = input port)  | Total count | Natural number |

| Category | Name             | Metric                              | Description  | Unit        | Value Range    |
|----------|------------------|-------------------------------------|--|-------------|----------------|
|          | PortXmitWait     | infiniband_port_transmit_wait_total | The number of ticks during which the port had data to transmit but no data was sent during the entire tick (either because of insufficient credits or because of lack of arbitration). | Total count | Natural number |
|          | PortXmitDiscards | infiniband_port_xmit_discards_total | Total number of outbound packets discarded by the port because the port is down or congested.  | Total count | Natural number |

**Table 4-6** Metric names

| Classification    | Metric            | Description   |
|-------------------|-------------------|---|
| Container metrics | modelarts_service | Service to which a container belongs, which can be <b>notebook</b> , <b>train</b> , or <b>infer</b>   |
|                   | instance_name     | Name of the pod to which the container belongs  |
|                   | service_id        | Instance or job ID displayed on the page, for example, <b>cf55829e-9bd3-48fa-8071-7ae870dae93a</b> for a development environment <b>9f322d5a-b1d2-4370-94df-5a87de27d36e</b> for a training job |
|                   | node_ip           | IP address of the node to which the container belongs   |
|                   | container_id      | Container ID  |
|                   | cid               | Cluster ID  |
|                   | container_name    | Name of the container   |
|                   | project_id        | Project ID of the account to which the user belongs   |
|                   | user_id           | User ID of the account to which the user who submits the job belongs  |

| Classification | Metric             | Description   |
|----------------|--------------------|---|
|                | pool_id            | ID of a resource pool corresponding to a physical dedicated resource pool             |
|                | pool_name          | Name of a resource pool corresponding to a physical dedicated resource pool           |
|                | logical_pool_id    | ID of a logical subpool   |
|                | logical_pool_name  | Name of a logical subpool   |
|                | gpu_uuid           | UUID of the GPU used by the container   |
|                | gpu_index          | Index of the GPU used by the container  |
|                | gpu_type           | Type of the GPU used by the container   |
|                | account_name       | Account name of the creator of a training, inference, or development environment task |
|                | user_name          | Username of the creator of a training, inference, or development environment task     |
|                | task_creation_time | Time when a training, inference, or development environment task is created           |
|                | task_name          | Name of a training, inference, or development environment task                        |
|                | task_spec_code     | Specifications of a training, inference, or development environment task              |
|                | cluster_name       | CCE cluster name  |
| Node metrics   | cid                | ID of the CCE cluster to which the node belongs                                       |
|                | node_ip            | IP address of the node  |
|                | host_name          | Hostname of a node  |
|                | pool_id            | ID of a resource pool corresponding to a physical dedicated resource pool             |
|                | project_id         | Project ID of the user in a physical dedicated resource pool                          |
|                | gpu_uuid           | UUID of a node GPU  |
|                | gpu_index          | Index of a node GPU   |

| Classification | Metric           | Description   |
|----------------|------------------|---|
|                | gpu_type         | Type of a node GPU  |
|                | device_name      | Device name of an InfiniBand or RoCE network NIC                          |
|                | port             | Port number of the InfiniBand NIC   |
|                | physical_state   | Status of each port on the InfiniBand NIC                                 |
|                | firmware_version | Firmware version of the InfiniBand NIC                                    |
|                | filesystem       | NFS-mounted file system   |
|                | mount_point      | NFS mount point   |
| Diagnos        | cid              | ID of the CCE cluster to which the node with the GPU equipped belongs     |
|                | node_ip          | IP address of the node where the GPU resides                              |
|                | pool_id          | ID of a resource pool corresponding to a physical dedicated resource pool |
|                | project_id       | Project ID of the user in a physical dedicated resource pool              |
|                | gpu_uuid         | GPU UUID  |
|                | gpu_index        | Index of a node GPU   |
|                | gpu_type         | Type of a node GPU  |
|                | device_name      | Name of a network device or disk device                                   |
|                | port             | Port number of the InfiniBand NIC   |
|                | physical_state   | Status of each port on the InfiniBand NIC                                 |
|                | firmware_version | Firmware version of the InfiniBand NIC                                    |