HPC Solution

User Guide

Issue 01

Date 2025-05-21

HUAWEI

HUAWEI TECHNOLOGIES CO., LTD.

Copyright © Huawei Technologies Co., Ltd. 2025. 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 Technologies Co., Ltd.

Trademarks and Permissions


and other Huawei trademarks are trademarks 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 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.

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process.* For details about this process, visit the following web page:
https://www.huawei.com/en/psirt/vul-response-process
For vulnerability information, enterprise customers can visit the following web page:
https://securitybulletin.huawei.com/enterprise/en/security-advisory

Contents

 Copyright © Huawei Technologies Co., Ltd.

 Copyright © Huawei Technologies Co., Ltd.

1 Overview

1.1 What's HPC-S²?

High-performance computing (HPC) is the practice of aggregating a large number of compute nodes to work concurrently to solve large-scale computing problems. It has been widely used in scientific research, weather forecast, simulation experiment, biopharmaceuticals, gene sequencing, and image processing.

The requirements for high-performance computing have been increasingly growing. HUAWEI CLOUD launched the HPC Solution Service (HPC-S²) to provide highly scalable high-performance computing resources as a cloud service. With basic services and software pre-integrated, HPC-S² allows you to provision and deploy HPC resources with just a few clicks. HPC-S² can automatically deploy job scheduling software such as Simple Linux Utility for Resource Management (Slurm), Sun Grid Engine (SGE), and Huawei commercial HPC Donau Scheduler to facilitate task assignments on parallel computers. On the HPC-S² console, you can view and manage clusters, nodes, and shared storage in an intuitive way.

You can also use the HPC-S² console to manage multiple clusters in different VPCs, as shown in **Figure 1-1**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 1-1 HPC-S² architecture

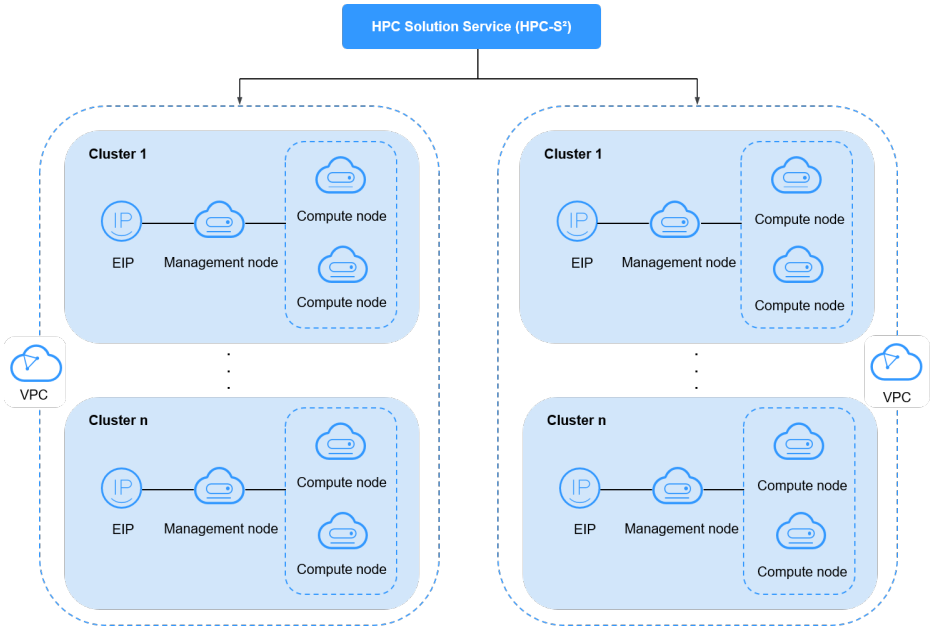

HPC-S² Functions

The HPC-S² is mainly used by enterprises that require a large amount of compute resources. These enterprises can use the HPC-S² console to easily manage clusters, nodes, tasks, and shared storage, provision clusters, and deploy applications.

Function	Description
Cluster management	Create HPC-S² clusters; view cluster creation, deployment, scale-in/out progress; view details about clusters; add, start, stop, and remove clusters; and manage instance topology.
Node management	Query, start, stop, restart, and delete nodes in HPC-S² clusters.
Shared storage management	Provide storage for HPC-S² clusters by using Scalable File Service (SFS).
Task management	Manage HPC-S² cluster-related job queueing, dynamic scheduling (including GPU scheduling), bin packing and workload consolidation, and quotas for CPU, memory, and instances.

HPC-S² Advantages

Compared with traditional HPC, HPC-S² offers the following advantages:

- HPC resources are provisioned on demand and are available instantly upon provisioning.

 Copyright © Huawei Technologies Co., Ltd.

- Multiple types of cloud servers are available for you to choose the most appropriate compute servers for your workloads.
- HPC-S² can work with security services, such as Anti-DDoS, Web Application Firewall (WAF), and Vulnerability Scan Service (VSS), to provide multi-dimensional protection to improve data security and availability.
- With HPC-S², you do not need to build your own equipment rooms or worry about hardware upgrade or replacement. You can obtain resources for exclusive use from virtual resource pools anytime and these resources can be elastically scaled based on service changes.
- HPC-S² can automatically configure password-free login, network information services (NISs), and schedulers for users.

1.2 HPC-S² Application and Usage Restrictions

HPC-S² Application Scenarios

- HPC-S² allows you to create HPC clusters with a few clicks. During the cluster creation, the schedulers, domain controllers, and runtime dependency are automatically deployed.
- You can manage clusters and nodes and view their status on the HPC-S² console.

HPC-S² Usage Restrictions

- The HPC-S² solution is available only in the AP-Singapore region.
- Only Slurm and SGE schedulers can be used for scheduling.
- The SGE scheduler can be in the Kunpeng or x86 architecture, but the Slurm scheduler can only be in the x86 architecture.
- A single user can create a maximum of three clusters.
- A cluster can contain a maximum of 2,000 compute nodes.
- Data disks cannot be automatically attached to compute nodes. If needed, you can manually attach them for use.
- If nodes fail to be created or added, the cluster cannot be automatically restored.
- Only CentOS 7 series public images can be used during cluster creation or capacity expansion.
- If the EIPs or data disks of the nodes are not deleted when the nodes are deleted, they will still incur fees.
- Do not perform operations on the nodes in the HPC cluster on the ECS console. Otherwise, the cluster may encounter with errors.
- Do not modify or delete files in the **/usr/local/.hpccluster** directory. Otherwise, the cluster may encounter errors.

  The **authorized_keys** file in the **/usr/local/.hpccluster/role/.ssh/** directory is used only for system maintenance.
- The clusters created by users with different member accounts cannot have duplicated names if such users belong to the same master account.

 Copyright © Huawei Technologies Co., Ltd.

1.3 Interaction Between HPC-S² and Other Services

Service	Interaction with HPC-S²	Reference
Elastic Cloud Server (ECS)	You can purchase ECSs to create HPC clusters or use existing ECSs to deploy HPC clusters.	**Purchasing a Custom ECS**
Virtual Private Cloud (VPC)	The cloud servers required for the same high-performance computing task are located in a single VPC, where subnets and security groups are configured for access control.	• **Creating a VPC and a Subnet** • **Modifying Details About a VPC**
Image Management Service (IMS)	You can use public, private, or shared images to create cloud servers required for high-performance computing. You can also create private images from the cloud servers fulfilling high-performance computing.	**Creating a Private Image**
Elastic Volume Service (EVS)	EVS disks can be attached to cloud servers to provide additional storage.	**Purchasing an EVS Disk**
Scalable File Service (SFS)	SFS provides hosted shared file storage for cloud servers.	**Sharing Files**

Copyright © Huawei Technologies Co., Ltd.

2 Getting Started

2.1 Overview

Scenarios

The typical HPC applications include DNA sequencing, industrial simulation, and other scenarios that require high performance computing.

This section describes how to create and configure an HPC cluster and how to complete HPC jobs on the HPC-S² console.

Methods of Creating an HPC Cluster

- **Purchasing an HPC cluster**: Purchase cloud servers with specified specifications as the master and compute nodes.

  Figure 2-1 Purchasing an HPC cluster

  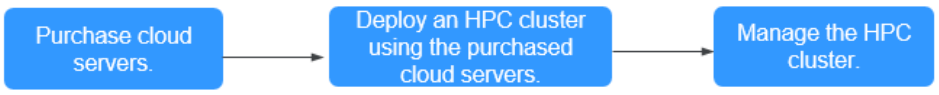

- **Deploying an HPC cluster**: Select existing cloud servers as the master and compute nodes.

  Figure 2-2 Deploying an HPC cluster using existing cloud servers

  

 Copyright © Huawei Technologies Co., Ltd.

2.2 Method 1: Purchasing an HPC Cluster

Scenarios

If you want to use new cloud servers to deploy an HPC cluster, perform operations in this section.

Notes

- When you use VPC for the first time, the system automatically creates a VPC for you, including a default security group and NIC. You can also create a VPC by yourself. For more information about VPC, see **Virtual Private Cloud User Guide**.
- You can use the default security group provided by the system or create a security group by yourself.

  To allow secure access to cloud servers in the cluster, ensure that port 22 in the security group rule is open. For details, see **Adding a Security Group Rule**.
- Scalable File Service (SFS) is used for shared file access spanning cloud servers. For more information about SFS, see **Scalable File Service User Guide**.

Procedure

1. Log in to the HPC-S^2 console.
2. Click  in the upper left corner and select your region and project.
3. Click **Buy HPC Cluster**.

   The **Buy HPC Cluster** page is displayed.
4. Select a billing mode, **Yearly/Monthly** or **Pay-per-use**.
   - **Yearly/Monthly**: Specify a required duration. The system will deduct the fees incurred from your account based on the service price.

     NOTE

     Yearly/Monthly cloud servers cannot be deleted. You can only unsubscribe from them. If such a cloud server is no longer needed, go to the corresponding service console to unsubscribe from it.
   - **Pay-per-use**: You do not need to set a required duration because the system will deduct the fees incurred from your account based on the service duration.
5. Select a region.

   Cloud servers in different regions cannot communicate with each other over an intranet. For low network latency and quick resource access, select the nearest region.
6. Select an AZ.

   An AZ is a physical location that uses an independent power supply and network. AZs in the same region can communicate with each other over an intranet.

 Copyright © Huawei Technologies Co., Ltd.

- To enhance high availability of applications, create cloud servers in different AZs.
- To shorten network latency, create cloud servers in the same AZ.

7. Configure the master node.
   - **Flavor**: Select a flavor from the drop-down list.
   - **System Disk**: Select a disk type and set the disk size.

     NOTE

     Before selecting a cloud server flavor, familiarize yourself with the cloud server types. For details, see **ECS Types**.
8. Configure compute nodes.
   - **Flavor**: Select a flavor from the drop-down list.
   - **System Disk**: Select a disk type and set the disk size.

     NOTE

     Before selecting a cloud server flavor, familiarize yourself with the cloud server types. For details, see **ECS Types**.
   - **Quantity**: You can deploy a maximum of 50 compute nodes in an HPC cluster.
9. Configure the network, including **VPC**, **Security Group**, and **NIC** settings.

   NOTE

   - The EIP is bound to the master node.
   - When you use VPC for the first time, the system automatically creates a VPC for you, including a default security group and NIC.

Table 2-1 Network parameters

Parameter	Description
VPC	A VPC provides a network, including subnets and security groups, for the cloud servers. Select an existing VPC or create one. For more information about VPC, see **Virtual Private Cloud User Guide**. **NOTE** DHCP must be enabled in the VPC to which the cloud servers belong.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Security Group	A security group controls access to the cloud servers in the group. By default, cloud servers in the same security group can access each other. The access from outside the security group is restricted by security group rules. You can define access rules for a security group to protect the cloud servers that are added to this group. For details, see **Adding a Security Group Rule**. Enable protocols and ports as needed. Common protocols and ports are as follows: • Port 80: default port for web page access through HTTP. • Port 443: web page access through HTTPS. • ICMP: checks communication statuses between cloud servers through ping operations. • Port 22: reserved for logging in to a Linux cloud server using SSH. • Port 3389: reserved for logging in to a Windows cloud server using SSH.
NIC	Select a subnet from the drop-down list.
EIP	An EIP is a static public IP address bound to a cloud server in a VPC. Using the EIP, the cloud server provides services externally. The following options are provided: • **Auto assign**: The system automatically assigns an EIP with a dedicated bandwidth to the cloud server. The bandwidth is configurable. • **Use existing**: An existing EIP will be assigned to the cloud server. The EIP is only bound to the master node.
EIP Type	Dynamic BGP provides automatic failover and chooses the optimal path when a network connection fails.
Bandwidth	This parameter is available only if **EIP** is set to **Auto assign**.

Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Billed By	This parameter is available only if **EIP** is set to **Auto assign**. This parameter indicates the bandwidth billing mode of the purchased EIP. Two billing modes are as follows: • **Bandwidth**: You will be billed based on the bandwidth size you set. • **Traffic**: You will be billed based on the total traffic used.

10. Select an image.
    - **Public image** (recommended)

      A public image is a standard, widely used image. It contains an OS and preinstalled public applications and is available to all users. If a public image does not contain the application environments or software you need, you can use the public image to create a cloud server and then deploy required software as needed.
    - **Private image**

      A private image is an image that contains an OS or service data, pre-installed public applications, and the user's personal applications. Private images are only available to the users who created them. Using a highly-tailored private image to create cloud servers eliminates the need to manually configure multiple cloud servers repeatedly.
    - **Shared image**

      A shared image is a private image another user has shared with you.
    - **Marketplace image**

      The Marketplace is an online store where you can purchase third-party images that have the OS, application environments, and software preinstalled. You can use these images to deploy websites and application development environments in just a few clicks. No additional configuration is required.
11. Set **Login Mode**.

    Key pair is recommended because it features higher security than **Password**. If you select **Password**, ensure that the password meets complexity requirements to prevent malicious attacks.
    - **Key pair**

      A key pair will be used for cloud server login authentication. You can select an existing key pair or create one.

      NOTE

      If you use an existing key pair, make sure that you have saved the key file locally. Otherwise, logging in to the cloud servers will fail.
    - **Password**

      A username and its initial password will be used for cloud server login authentication.

 Copyright © Huawei Technologies Co., Ltd.

The initial password of user **root** will be used for authenticating Linux login, while that of user **Administrator** will be used for authenticating Windows login. The password must meet the following requirements.

Table 2-2 Password complexity requirements

Item	Requirements	Example Value
Password	• Consists of 8 to 26 characters. • Contains at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters: $!@%-_=+[]:./^, {}? • Cannot contain the username or the username spelled backwards. • Cannot contain more than two characters in the same sequence as they appear in the username. (This requirement applies only to Windows cloud servers.)	YNbUwp! dUc9MClnv **NOTE** The example password is generated randomly. Do not copy it for use.

NOTE

The system does not periodically change the password. It is recommended that you change it regularly for security.

12. Set a cluster name.

    The name must comply with the naming rules in the following table.

Table 2-3 Cluster naming rules

Item	Requirements	Example Value
Cluster name	• Consists of 1 to 52 characters. • Contains only letters, digits, hyphens (-), and periods (.). • Cannot start or end with a hyphen (-) or period (.). • Cannot contain consecutive hyphens (-) and periods (.) or the combination of them, for example, **--**, **..**, **-.**, or **.-**.	hpc-001.p1

13. Select a file system.

 Copyright © Huawei Technologies Co., Ltd.

Available file systems are filtered and displayed in the drop-down list based on the network settings. If no files systems are available, you can create one and add it to the VPC specified in **Table 2-1**. For details, see **Create a File System**.

14. Select a scheduler.

    Currently, only the **slurm** and **sge** schedulers are available. The default one is **slurm**.

15. Confirm the configurations and click **Buy Now**. A page is displayed for you to confirm the order.
    - In **Yearly/Monthly** billing mode, after confirming the specifications and price of the order, agree to the service agreement, click **Submit**, and complete the payment as prompted.
    - In **Pay-per-use** billing mode, after confirming the specifications and price of the order, agree to the service agreement, and click **Submit**.
16. After the order is submitted or paid, the system returns back to the **Clusters** page, on which you can see that the cluster is in the **Creating** state. Wait until the cluster status changes to **Running** which indicates that the cluster is successfully created.

Follow-up Operations

View the cluster details, nodes, and shared storage, as instructed in **Viewing HPC Clusters**.

2.3 Method 2: Deploying an HPC Cluster

Scenarios

If you want to use existing cloud servers to deploy an HPC cluster, perform operations in this section.

NOTICE

Select a master node and then compute nodes. These nodes must have the same CPU architecture and OS version and in the same VPC and security groups.

Notes

- When you use VPC for the first time, the system automatically creates a VPC for you, including a default security group and NIC. You can also create a VPC by yourself. For more information about VPC, see **Virtual Private Cloud User Guide**.
- You can use the default security group provided by the system or create a security group by yourself.

  To allow secure access to cloud servers in the cluster, ensure that port 22 in the security group rule is open. For details, see **Adding a Security Group Rule**.

 Copyright © Huawei Technologies Co., Ltd.

- Scalable File Service (SFS) is used for shared file access spanning cloud servers. For more information about SFS, see **Scalable File Service User Guide**.

Procedure

1. Log in to the HPC-S^2 console.
2. Click  in the upper left corner and select your region and project.
3. Click **Deploy HPC Cluster**.

   The **Deploy HPC Cluster** page is displayed.
4. Select a region.

   Cloud servers in different regions cannot communicate with each other over an intranet. For low network latency and quick resource access, select the nearest region.
5. Set a cluster name.

   The name must comply with the naming rules in the following table.

   Table 2-4 Cluster naming rules

Item	Requirements	Example Value
Cluster name	• Consists of 1 to 52 characters. • Contains only letters, digits, hyphens (-), and periods (.). • Cannot start or end with a hyphen (-) or period (.). • Cannot contain consecutive hyphens (-) and periods (.) or the combination of them, for example, **--**, **..**, **-.**, or **.-**.	hpc-001.p1

6. Select a cloud server as the master node.

   The master node must:
   - Be in the region selected in **2**.
   - Be in the **Running** state.
   - Have an EIP bound.
   - Not be included in other clusters under the current account.
7. Select security groups.

   A security group controls the access to the cloud servers added to this group by defining access rules. You can select multiple security groups.
8. Select compute nodes.

   The compute nodes must:
   - Be in the region selected in **2**.
   - Be in the **Running** state.
   - Use the same CPU architecture as the master node.

 Copyright © Huawei Technologies Co., Ltd.

- Be in the same VPC and security groups as the master node.
- Use the same major OS version as the master node.
- Be no more than 50 in quantity.

9. Select a VPC.

   This parameter is automatically set to the VPC of the master node.

10. Set the login mode.

    Only **Password** is available. If any of the selected cloud servers is using key pair authentication, go to the corresponding service console to change it to password authentication. In addition, set a strong password to protect yourself from malicious attacks.

    The initial password of user **root** will be used for authenticating Linux login, while that of user **Administrator** will be used for authenticating Windows login. The password must meet the following requirements.

    Table 2-5 Password complexity requirements

Item	Requirements	Example Value
Password	• Consists of 8 to 26 characters. • Contains at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters: $!@%-_=+[]:./^,{}? • Cannot contain the username or the username spelled backwards. • Cannot contain more than two characters in the same sequence as they appear in the username. (This requirement applies only to Windows cloud servers.)	YNbUwp!dUc9MClnv **NOTE** The example password is generated randomly. Do not copy it for use.

11. Select a file system.

    Available file systems are filtered and displayed in the drop-down list based on the network settings. If no files systems are available, you can create one and add it to the VPC specified in **9**. For details, see **Create a File System**.

12. Select a scheduler.

    Currently, only the **slurm** and **sge** schedulers are available. The default one is **slurm**.

13. Confirm the configurations and click **Deploy HPC Cluster**.

14. After the order is submitted, the system returns back to the **Clusters** page, where you can see that the cluster is in the **Deploying** state. Wait until the cluster status changes to **Running** which indicates that the cluster is successfully deployed.

 Copyright © Huawei Technologies Co., Ltd.

Follow-up Operations

View the cluster details, nodes, and shared storage, as instructed in **Viewing HPC Clusters**.

2.4 Getting Started with the HPC Management and Scheduling Plug-in

2.4.1 Deployment

Preparations

Provide a Huawei Cloud account, an AK/SK pair, and the region information.

Obtaining an Image

1. Request an image that contains the plug-in from Huawei after-sales personnel.
2. (Huawei Cloud after-sales personnel) Share the image.
3. Log in to the Huawei Cloud management console and choose **Compute** > **Image Management Service**. Then, receive the image and copy it to a private one.

Purchasing an SFS Turbo System

The cockpit system depends on SFS Turbo, so purchase an SFS Turbo file system first.

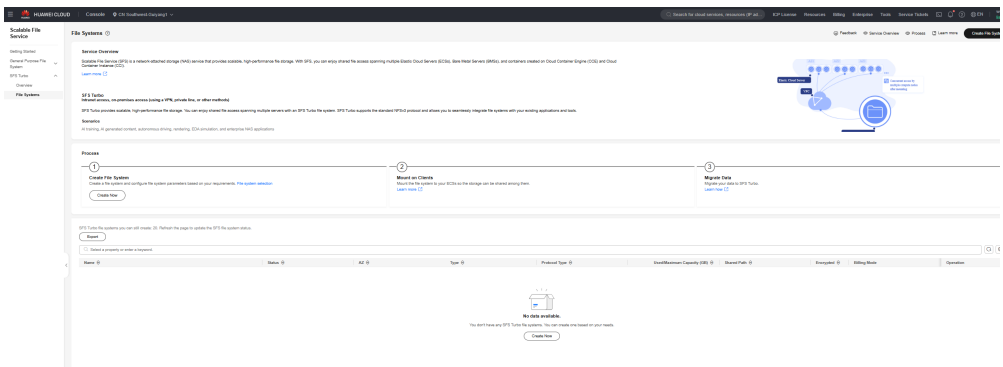

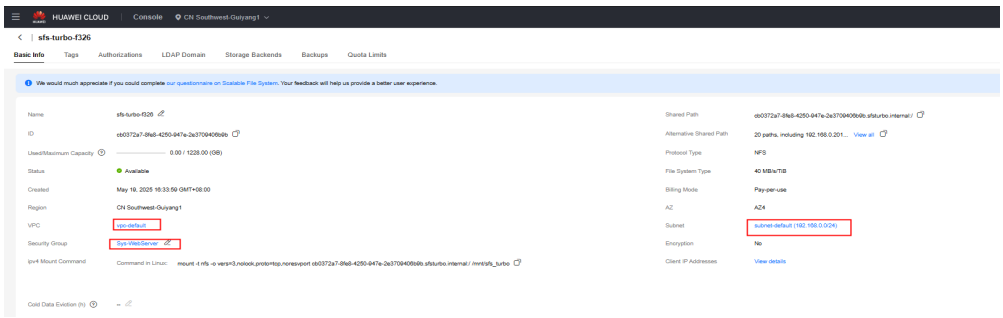

 Copyright © Huawei Technologies Co., Ltd.

Cockpit Master Node Deployment

Select the private image of the master node. Use the security group, subnet, and VPC that are same as the SFS Turbo file system. Choose the flavor with 16 vCPUs and 32 GB memory, set the system disk to 300 GB, and bind an EIP.

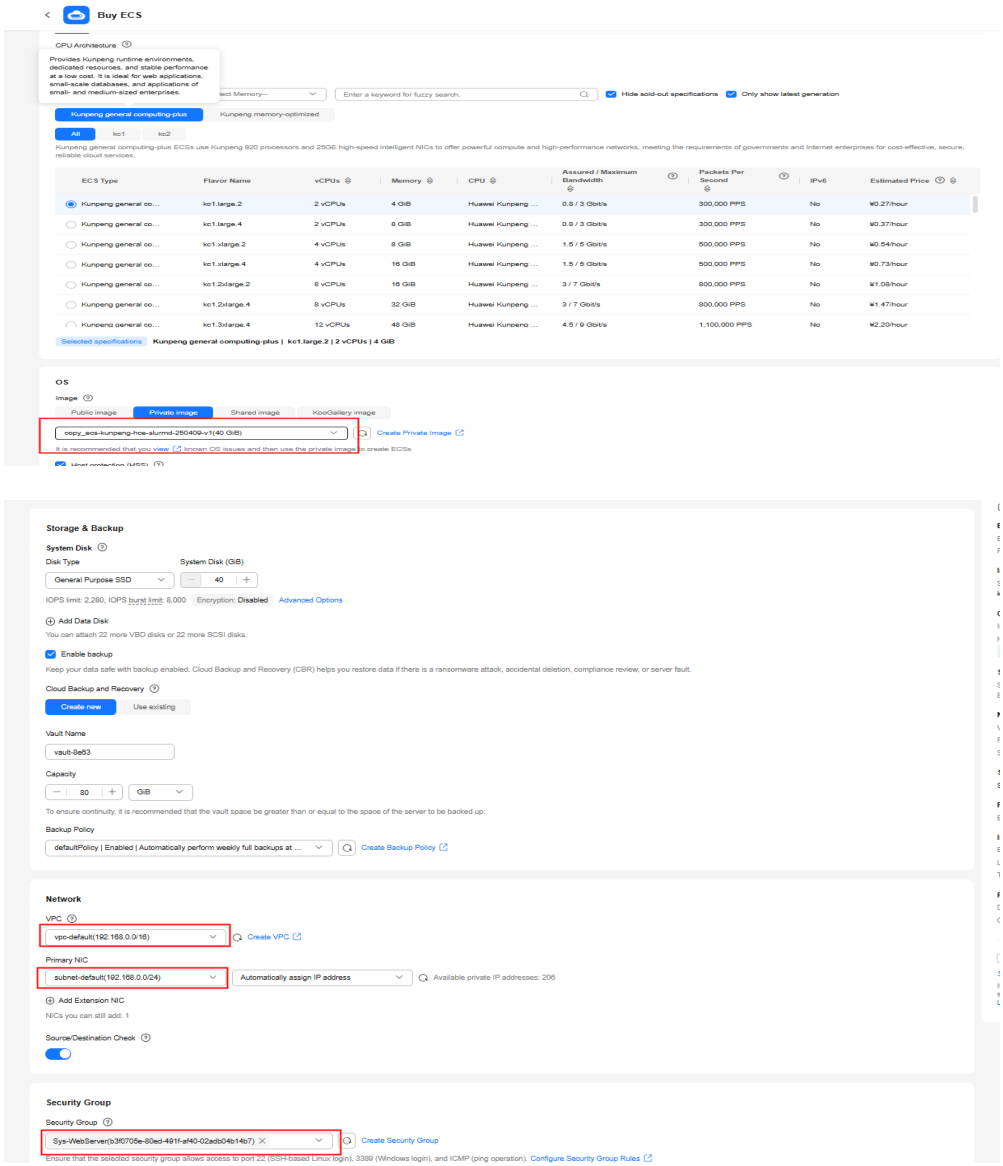

Under **Advanced Settings**, enter the following custom script, set **SFS_ENDPOINT** to an actual value, and click **Buy**.

```
#!/bin/bash
cat >>/opt/cloud/services/intelligent-cockpit-agent/user_data.sh  <<EOF
CLUSTER_NAME=IntelligentCockpit
DEPLOYMENT_MODE=SIMPLE
USER_MANAGER_MODE=NIS
NIS_DOMAIN=intelligentCockpitServer.com
SFS_ENDPOINT=xxx
EOF

systemctl daemon-reload
systemctl start  CockpitStart.service
systemctl enable  CockpitStart.service
```

 Copyright © Huawei Technologies Co., Ltd.

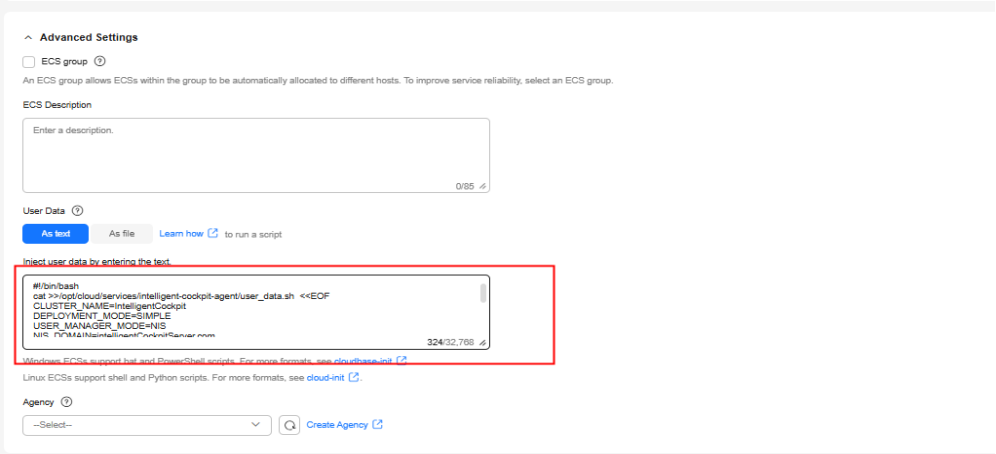

After the purchase is successful, configure the variables in the following script and set **master-ip** to the private IP address of the master node. Then, log in to the master node and execute the script. After the script is executed, the master node of the cockpit has been deployed.

```
#!/bin/bash
cat >> /opt/cloud/services/intelligent-cockpit-agent/application.properties  <<EOF
cluster.master-ip=xxx
cluster.intelligent-cockpit-master-ip=xxx
cluster.slurm-rest-url=http://127.0.0.1:6688
cluster.regionId=cn-southwest-2
user.domainId=xxx
user.ak=xxx
user.sk=xxx
user.project-id=xxx
cluster.sfs-share-ids=xxx
cluster.nis-domain-name=intelligentCockpitServer.com
EOF

systemctl daemon-reload
systemctl start  IntelligentCockpitAgent.service
systemctl enable  IntelligentCockpitAgent.service
```

2.4.2 Quick Start

1. Log in to HPC.

   Visit the EIP associated with the master node of the cockpit to switch to the login page. Then, use **admin/admin** for login.

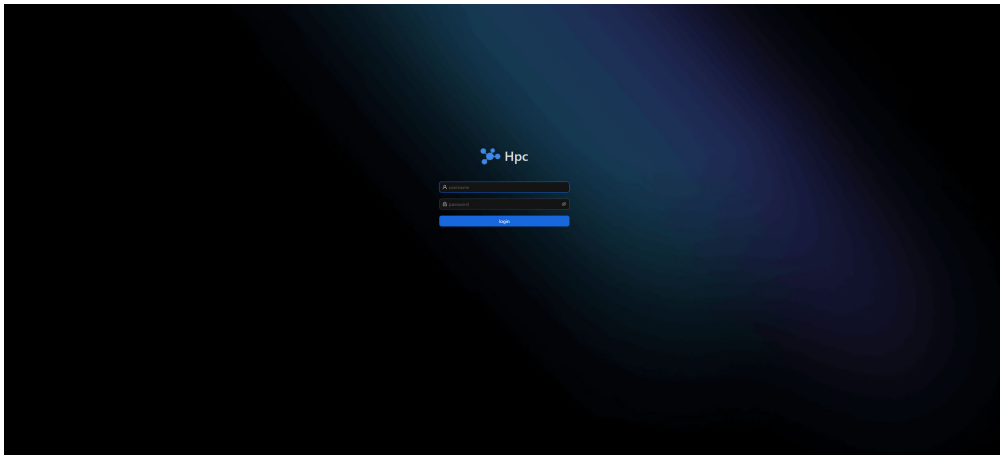

 Copyright © Huawei Technologies Co., Ltd.

2. Add a partition.

   On the partition management page, choose **Cluster Management** > **Partition Management**.

   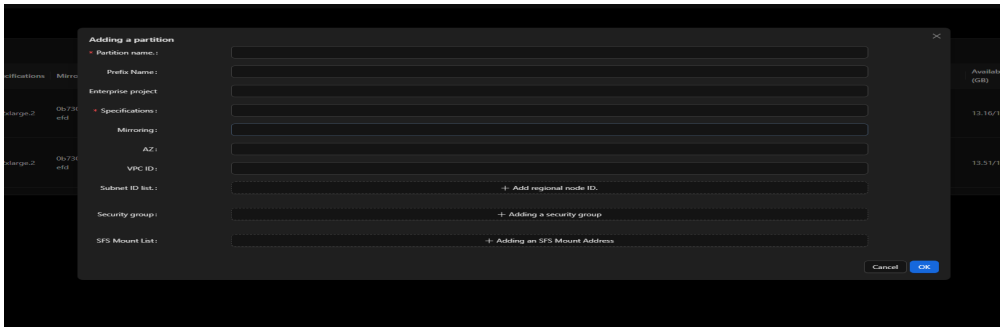

   The parameters are described as follows.

Parameter	Description
Specifications	Compute node flavor, for example, **c6s.xlarge.2**
Mirroring	ID of the private image of the compute nodes
AZ	AZ where compute nodes are located, for example, **cn-southwest-2a**
VPC ID	ID of the VPC where compute nodes are located. It is recommended that the VPC ID be the same as that of the master node.
Subnet ID list	ID of the subnet where compute nodes are located. It is recommended that the subnet ID be the same as that of the master node.
Security group	ID of the security group where compute nodes are located. It is recommended that the security group ID be the same as that of the master node.
SFS Mount List	Mount address of the SFS Turbo file system. For example, if *xxx***.sfsturbo.internal:/** is the sharing path, you need to mount *xxx***.sfsturbo.internal:/** to **/mnt/sfs_turbo_1** and mount *xxx***.sfsturbo.internal:/home** to **/home**.

3. Create a compute cluster.

   Go to the **Resource Management** page and click **Add Compute Node**.

 Copyright © Huawei Technologies Co., Ltd.

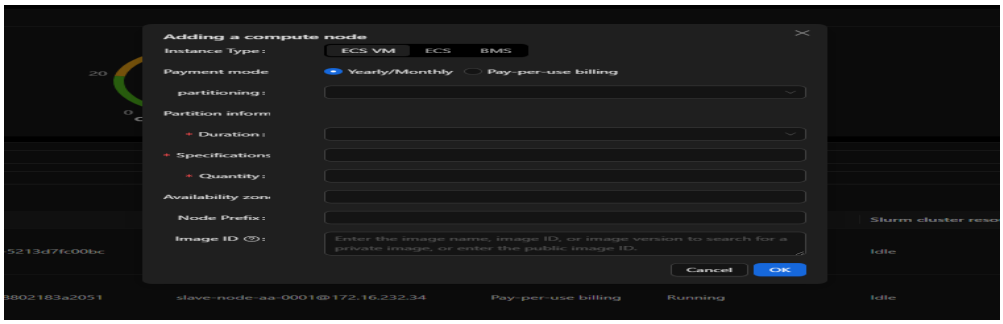

The parameters are described as follows.

Parameter	Description
Payment mode	Billing mode of compute nodes, which can be **Yearly/Monthly** or **Pay-per-use**.
Partitioning	Partition where nodes are located
Specifications	Compute node flavor
Quantity	Number of compute nodes to create
Availability zone	(Optional) AZ where the compute nodes are located. If this parameter is not set, the AZ set for the partition will be used.
Node Prefix	(Optional) Compute node prefix. If this parameter is not set, the prefix set for the partition will be used.
Image ID	(Optional) Image used by compute nodes. If this parameter is not set, the image set for the partition will be used.

4. Create a template.

   On the template management page, choose **System Management** > **Job Template Management**.

5. Submit a job.

   On the service management page, submit a job.

 Copyright © Huawei Technologies Co., Ltd.

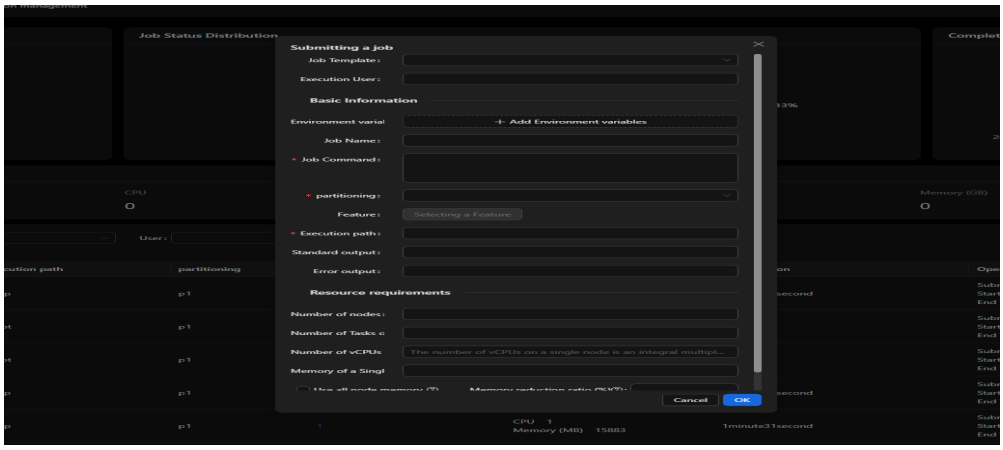

The parameters are described as follows.

Parameter	Description
Job Template	Choose a template from the drop-down list.
Execution User	User who executes tasks. By default, the current user is used.
Environment variables	Environment variables that need to be transferred to Slurm
Job Command	Task script
partitioning	Partition to which tasks are submitted
Feature	Features can be added to compute nodes. When tasks are submitted, they can be scheduled based on the features.
Execution path	Path where the script is executed. The path must be a mounted public path.
Standard output	(Optional) Name of the log file generated by the job
Error output	(Optional) Exception log file of a job
Number of nodes	Number of nodes required by a job
Number of Tasks on a Single Node	Number of tasks running on each node
Number of vCPUs on a Single Node	Number of vCPUs on each node
Memory of a Single Node	Single-node memory required by tasks

 Copyright © Huawei Technologies Co., Ltd.

After a job is submitted, you can check job logs in real time.

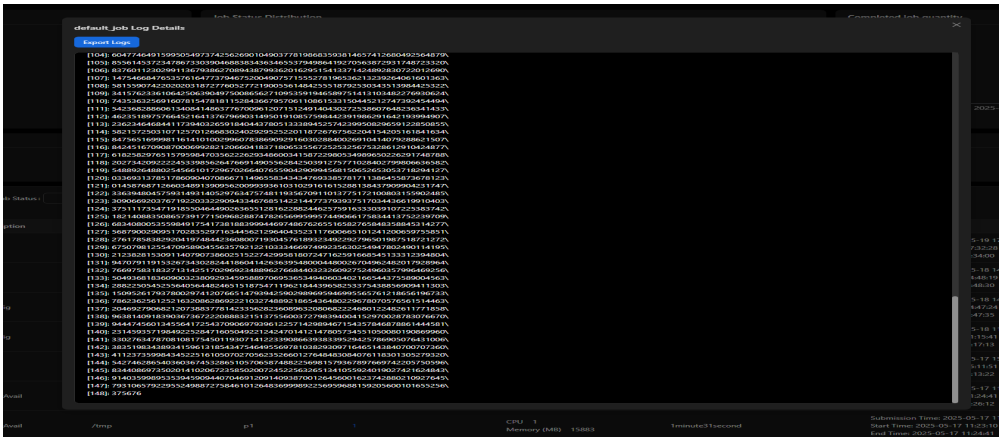

After the job is complete, you can check the job list, statistics, and resource usage.

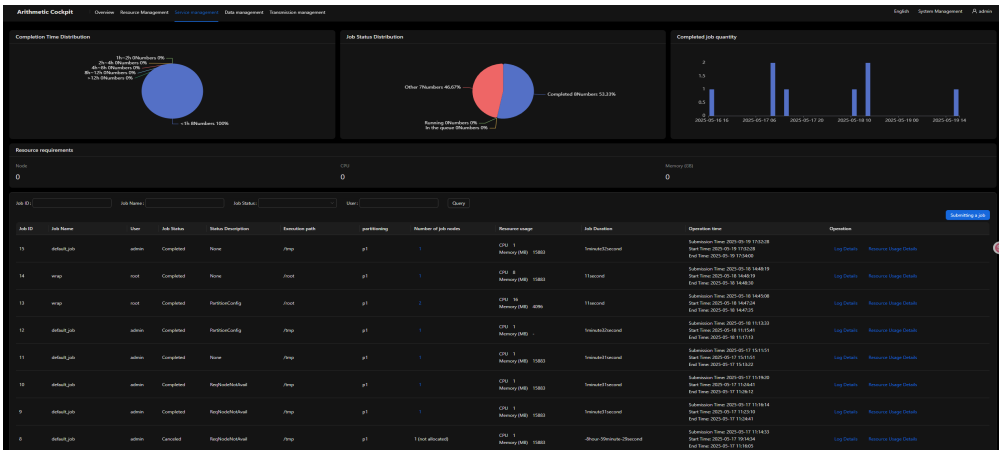

2.4.3 Core Features

Data Management

Files smaller than 1 GB can be quickly uploaded or downloaded in the cockpit. For files larger than 1 GB, you are advised to upload and download them using OBS to avoid bandwidth or browser restrictions.

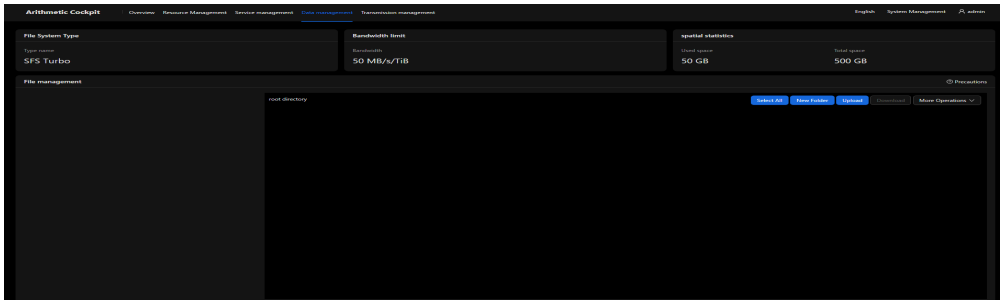

 Copyright © Huawei Technologies Co., Ltd.

Elastic Resource Provisioning

The cockpit allows you to associate an auto scaling policy with each partition.

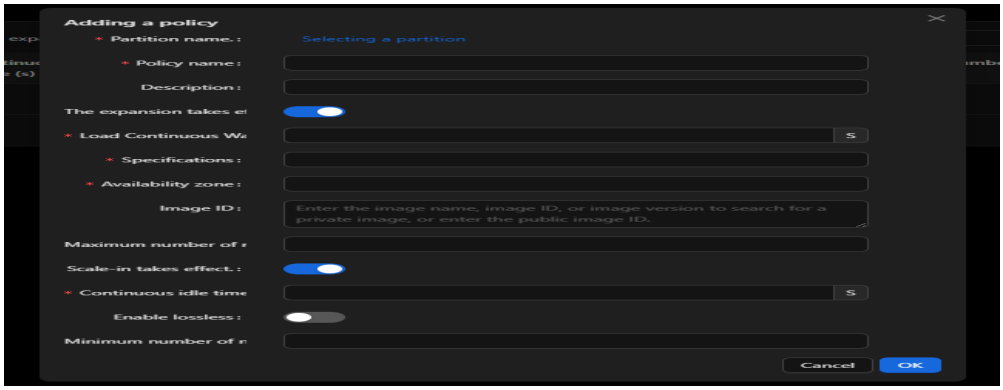

The parameters are described as follows.

Parameter	Description
Partition name	Partition which an auto scaling policy is associated with. Each partition can have only one policy.
Policy name	Name of the auto scaling policy
Load Continuous Waiting Time	If scaling out is enabled and the continuous load wait time is longer than the specified time, the load is considered valid. Compute nodes then are calculated and added.
Specifications	Flavor of the compute nodes to be added, for example, **c6s.xlarge.2**
Availability zone	AZ where the compute nodes to be added are located, for example, **cn-southwest-2a**
Image ID	ID of the private image used by the compute nodes to be added
Maximum number of nodes	Maximum number of compute nodes in a partition
Continuous idle time	Continuous idle time of a node. If scaling in is enabled and the nodes are idle for a period longer than the specified duration, the nodes will be deleted.
Minimum number of nodes	Minimum number of nodes that should be reserved in a partition. If the left nodes are fewer than the specified number, scaling in will not be performed.

 Copyright © Huawei Technologies Co., Ltd.

Tag Management

You can add tags to nodes in a partition. When tasks are submitted, they can be scheduled based on the tags.

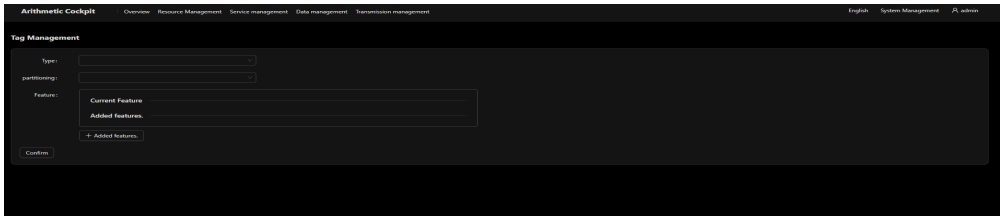

 Copyright © Huawei Technologies Co., Ltd.

3 User Guide

3.1 Viewing the Cluster Deployment Progress

3.1.1 Viewing the Progress of Purchasing an HPC Cluster

Scenarios

After you submit the order for purchasing an HPC cluster, you can view the cluster deployment progress on the HPC-S² console.

Procedure

1. Log in to the HPC-S² console.
2. In the navigation pane on the left, choose **Clusters**.
3. In the cluster list, locate the row where the cluster you want to view resides and click **View Progress** in the **Operation** column.

   On the displayed page, set **Log Type** to **Deployment** and view the deployment progress of the cluster.

3.1.2 Viewing the Progress of Deploying an HPC Cluster Using Existing Cloud Servers

Scenarios

After you submit a request for deploying an HPC cluster using existing cloud servers, you can view the cluster deployment progress on the HPC-S² console.

Procedure

1. Log in to the HPC-S² console.
2. In the navigation pane on the left, choose **Clusters**.

 Copyright © Huawei Technologies Co., Ltd.

3. In the cluster list, locate the row where the cluster you want to view resides and click **View Progress** in the **Operation** column. On the displayed page, set **Log Type** to **Deployment** and view the deployment progress of the cluster.

3.1.3 Viewing the Scale-Out Progress of an HPC Cluster

Scenarios

After you submit a request for scaling out an HPC cluster, you can view the scale-out progress on the HPC-S^2 console.

Procedure

1. Log in to the HPC-S^2 console.
2. In the navigation pane on the left, choose **Clusters**.
3. In the cluster list, locate the row where the cluster you want to view resides and click **View Progress** in the **Operation** column.

   On the displayed page, set **Log Type** to **Manual scale-out** and view the scale-out progress.

3.1.4 Viewing the Scale-In Progress of an HPC Cluster

Scenarios

After you submit a request for scaling in an HPC cluster, you can view the scale-in progress on the HPC-S^2 console.

Procedure

1. Log in to the HPC-S^2 console.
2. In the navigation pane on the left, choose **Clusters**.
3. In the cluster list, locate the row where the cluster you want to view resides and click **View Progress** in the **Operation** column.

   On the displayed page, set **Log Type** to **Manual scale-in** and view the scale-in progress.

3.2 Viewing HPC Clusters

3.2.1 Viewing Details of an HPC Cluster

Scenarios

You can view details about an HPC cluster, including the cluster ID, status, region, VPC, security groups, nodes, shared storage, and topology.

Procedure

1. Log in to the HPC-S^2 console.

 Copyright © Huawei Technologies Co., Ltd.

2. In the navigation pane on the left, choose **Clusters**.
3. In the cluster list, click the name of the cluster you want to view.

   On the cluster details page, you can view the cluster ID, status, region, VPC, security groups, nodes, shared storage, and topology.

3.2.2 Viewing Details About Nodes in an HPC Cluster

Scenarios

You can view node details of an HPC cluster, including the node name/ID, type, status, specifications, IP address, shared storage, AZ, and billing mode.

Procedure

1. Log in to the HPC-S² console.
2. In the navigation pane on the left, choose **Nodes**.
3. On the displayed page, select the cluster you want to view from the cluster drop-down list. The nodes in the cluster are displayed.
4. Locate the node you want to view and click its name to go to the ECS console for its details.

3.2.3 Managing Shared Storage of an HPC Cluster

Scenarios

You can view details about the shared storage used by an HPC cluster, including the name and ID of the SFS file system, protocol, shared path, and mount point.

Procedure

1. Log in to the HPC-S² console.
2. In the navigation pane on the left, choose **Shared Storage**.
3. On the displayed page, select the cluster you want to view from the cluster drop-down list, and view details about the shared storage of the cluster.

3.3 HPC Cluster Management

3.3.1 Scaling Out an HPC Cluster

Scenarios

If the HPC cluster does not meet your requirements, you can scale out the cluster.

NOTE

Only clusters in the **Running** state can be scaled out.

 Copyright © Huawei Technologies Co., Ltd.

Procedure

1. Log in to the HPC-S^2 console.
2. In the navigation pane on the left, choose **Clusters**.
3. On the displayed page, locate the row where the cluster you want to scale out resides and click **Scale** in the **Operation** column. The scale-out details page is displayed.
4. Select a billing mode.

   NOTE

   During scale-out, you can select a different billing mode for scaled nodes. If the cluster contains both the pay-per-use nodes and yearly/monthly billing nodes, the cluster billing mode is displayed as **Yearly/Monthly and Pay-per-use**.

   - **Yearly/Monthly**: Specify a required duration. The system will deduct the fees incurred from your account based on the service price.
   - **Pay-per-use**: You do not need to set a required duration because the system will deduct the fees incurred from your account based on the service duration.
5. Specify **Region** and **AZ**.
   - By default, the region is unconfigurable and set to the region where the cluster is created.
   - You can select a different AZ.
6. Configure the specifications for new nodes.
   - **Flavor**: Select a flavor from the drop-down list.
   - **System Disk**: Select a disk type and set the disk size.
7. Configure the network.

   VPC, **Security Group**, and **NIC** are automatically set to those configured when the cluster was created.
8. Select an image.

   NOTE

   You are advised to select a public image for new nodes. Otherwise, the scale-out may fail.

   - **Public image** (recommended)

     A public image is a standard, widely used image that is visible to all users. It contains an OS and preinstalled public applications. Currently, only CentOS 7 images are available.
   - **Private image**

     A private image is only available to the user who created it. It contains an OS, preinstalled public applications, and the user's personal applications. Using a highly-tailored private image to create cloud servers eliminates the need to manually configure multiple cloud servers repeatedly.

     You can also select an encrypted image. For details, see *Image Management Service User Guide*.
   - **Shared image**

     A shared image is a private image another user has shared with you.

 Copyright © Huawei Technologies Co., Ltd.

- **Marketplace image**

  The Marketplace is an online store where you can purchase third-party images that have the OS, application environments, and software preinstalled.

9. Set **Login Mode**.

   NOTE

   The login mode must be the same as that configured when the cluster was created.

   - **Key pair**

     A key pair will be used for cloud server login authentication. You can select an existing key pair or create one.

     NOTE

     If you use an existing key pair, make sure that you have saved the key file locally. Otherwise, logging in to the cloud servers will fail.

   - **Password**

     A username and its initial password will be used for cloud server login authentication.

     The password must meet the following requirements.

Item	Requirement	Example Value
Password	• Consists of 8 to 26 characters. • Contains at least three of the following character types: • Uppercase letters • Lowercase letters • Digits • Special characters: $!@%-_=+[]:./^,{}? • Cannot contain the username or the username spelled backwards.	YNbUwp!dUc9MClnv **NOTE** The example password is generated randomly. Do not copy it for use.

10. Select the amount of time you want to use the nodes for.

    If you set **Billing Mode** to **Yearly/Monthly**, you need to select the required duration and determine whether to select **Auto-renew**. If you select **Auto-renew**, the node will be automatically renewed when the subscription period expires.

11. Specify **Quantity**.

    The total number of existing and new compute nodes in the cluster cannot exceed 50.

12. Confirm the configurations and click **Buy Now**. A page is displayed for you to confirm the order.

 Copyright © Huawei Technologies Co., Ltd.

- In **Yearly/Monthly** billing mode, after confirming the specifications and price of the order, agree to the service agreement, click **Submit**, and complete the payment as prompted.
- In **Pay-per-use** billing mode, after confirming the specifications and price of the order, agree to the service agreement, and click **Submit**.

13. After the order is submitted or paid, the system returns back to the **Clusters** page where you can see that the cluster is in the **Creating** state. Wait until the cluster status changes to **Running**.

3.3.2 Stopping an HPC Cluster

Scenarios

If there are no jobs to run on a pay-per-use HPC cluster, you can stop the cluster to save money.

NOTE

- Yearly/monthly cloud servers will continue to be billed after the cluster is stopped. You need to manually unsubscribe from the cloud servers to stop billing.
- Only clusters in the **Running** state can be stopped.

Procedure

1. Log in to the HPC-S² console.
2. In the navigation pane on the left, choose **Clusters**.
3. In the cluster list, locate the row where the cluster you want to stop resides, and choose **More** > **Stop** in the **Operation** column.
4. Click **OK**.
5. Wait until the cluster status changes to **Stopped**.

3.3.3 Starting an HPC Cluster

Scenarios

If you need to execute jobs on a stopped cluster, you can start the cluster.

NOTE

Only clusters in the **Stopped** state can be started.

Procedure

1. Log in to the HPC-S² console.
2. In the navigation pane on the left, choose **Clusters**.
3. In the cluster list, locate the row where the cluster you want to start resides, and choose **More** > **Start** in the **Operation** column.
4. Click **Next**.
5. Wait until the cluster status changes to **Running**.

 Copyright © Huawei Technologies Co., Ltd.

3.3.4 Deleting an HPC Cluster

Scenarios

- If a cluster is no longer needed, you can delete it. If you do not select **Delete cloud server** when you delete the cluster, these servers will be retained and continue to be billed after the cluster is deleted. To stop billing for these cloud servers, manually delete pay-per-use servers and unsubscribe from yearly/monthly servers.
- If a cluster fails to be created, you can delete the cluster and do not select **Delete cloud server**. Then, use these servers to deploy a cluster again.
- If an exception occurs in a cluster and the cluster cannot be manually restored, you can delete the cluster and do not select **Delete cloud server**. Then, use these servers to deploy a cluster again.

Precautions

- The EIP bound to the cluster will continue to be billed after the cluster is deleted. You need to manually delete the EIP to stop billing.
- Data disks attached to nodes in the cluster will continue to be billed after the cluster is deleted, even if you select **Delete cloud server** when deleting the cluster. You need to manually delete the data disks to stop billing.

Procedure

1. Log in to the HPC-S^2 console.
1. In the navigation pane on the left, choose **Clusters**.
2. In the cluster list, locate the row where the cluster you want to delete, and choose **More** > **Delete** in the **Operation** column. On the confirmation page, whether option **Delete cloud server** is available depends on the billing mode of the cluster you want to delete.
   - Pay-per-use

     Delete cloud server is available on the confirmation page.

     - Option selected

       The cloud servers will be deleted together with the cluster.

     - Option not selected

       The cloud servers are retained. You can view the cloud servers on the corresponding service console.
   - Yearly/Monthly

     Delete cloud server is unavailable on the confirmation page.

     You need to unsubscribe from the yearly/monthly cloud servers in the cluster.
   - Yearly/Monthly and Pay-per-use

     Delete cloud server is available on the confirmation page.

     - Option selected

 Copyright © Huawei Technologies Co., Ltd.

Pay-per-use cloud servers will be deleted together with the cluster. You need to manually unsubscribe from the yearly/monthly cloud servers in the cluster.

- Option not selected

  The cloud servers are retained. You can view the cloud servers on the corresponding service console.

3.4 Node Management

3.4.1 Stopping a Node

Scenarios

If a pay-per-use node is faulty or no longer needed, you can stop the node to save money.

NOTE

- A yearly/monthly node will continue to be billed after being stopped. You need to manually unsubscribe from the node to stop billing.
- Only nodes in the **Running** state can be stopped.

Procedure

1. Log in to the HPC-S² console.
2. In the navigation pane on the left, choose **Nodes**.
3. Select the target cluster from the cluster drop-down list.
4. Locate the row where the node you want to stop resides, and choose **More** > **Stop** in the **Operation** column.
5. Click **Yes**.
6. Wait until the node status changes to **Stopped**.

3.4.2 Starting a Node

Scenarios

You can start nodes as needed.

Only nodes in the **Stopped** state can be started.

Procedure

1. Log in to the HPC-S² console.
2. In the navigation pane on the left, choose **Nodes**.
3. Select the target cluster from the cluster drop-down list.
4. Locate the row where the node you want to start resides and choose **More** > **Start** in the **Operation** column.

 Copyright © Huawei Technologies Co., Ltd.

5. Click **Yes**.
6. Wait until the node status changes to **Running**.

3.4.3 Restarting a Node

Scenarios

You can restart a faulty node to restore it.

Only nodes in the **Running** state can be restarted.

Procedure

1. Log in to the HPC-S² console.
2. In the navigation pane on the left, choose **Nodes**.
3. Select the target cluster from the cluster drop-down list.
4. Locate the row where the node you want to restart resides and choose **More** > **Restart** in the **Operation** column.
5. Click **Yes**.
6. Wait until the node status changes to **Running**.

3.4.4 Removing a Node

Scenarios

- If there are too many idle nodes in the cluster, you can remove some nodes to save money.
- If a faulty node in the cluster cannot be manually restored, you can remove it and then add a new node to the cluster as needed.

Precautions

- Master nodes cannot be removed.
- A compute node can be removed only when its cluster is in the **Running** or **Abnormal** state.
- A compute node can be removed only when other nodes in the cluster are in the **Running** state.
- Before removing a compute node, you need to ensure that no jobs are running on the node. Otherwise, the running jobs will be abnormal.
- Attached data disks will be retained and continue to be billed after the node is removed. To stop billing for these data disks, delete them manually.

Removing a Single Node

1. Log in to the HPC-S² console.
2. In the navigation pane on the left, choose **Nodes**.
3. Select the target cluster from the cluster drop-down list.

   Locate the row where the node you want to remove resides and choose **More** > **Remove** in the **Operation** column.

 Copyright © Huawei Technologies Co., Ltd.

- Pay-per-use

  If you select **Delete cloud server**, the cloud server will be deleted. Otherwise, the cloud server will be retained. You can view the cloud server on the corresponding service console.

- Yearly/Monthly

  You need to manually unsubscribe from yearly/monthly cloud servers.

Removing Nodes in Batches

1. Log in to the HPC-S^2 console.
2. In the navigation pane on the left, choose **Nodes**.
3. Select the target cluster from the cluster drop-down list.

   Select the nodes to be removed and click **Remove** on the left of the cluster name.

   - Pay-per-use

     If you select **Delete cloud server**, the cloud servers will be deleted. Otherwise, the cloud servers will be retained. You can view the cloud servers on the corresponding service console.

   - Yearly/Monthly

     You need to manually unsubscribe from yearly/monthly cloud servers.

   - Yearly/Monthly and Pay-per-use

     If you select **Delete cloud server**, pay-per-use cloud servers will be deleted and yearly/monthly cloud servers will be retained and continue to be billed. To stop billing for these yearly/monthly cloud servers, unsubscribe from them manually. If you do not select **Delete cloud server**, all cloud servers will be retained.

3.5 Shared Storage Management

3.5.1 Managing Shared Storage

Scenarios

The HPC-S^2 console enables you to mount SFS file systems to and unmount them from running clusters. SFS provides shared storage for clusters.

Mounting

1. Log in to the HPC-S^2 console.
2. In the navigation pane on the left, choose **Shared Storage**.
3. On the **Shared Storage** page, select the target cluster from the drop-down list in the upper left corner.
4. In the upper right corner of the displayed page, click **Mount**.
5. On the **Mount Shared Storage** page, enter the mount point for the shared storage and select a file system from the drop-down list.

 Copyright © Huawei Technologies Co., Ltd.

6. Click **Yes**.

Unmounting

1. Log in to the HPC-S² console.
2. In the navigation pane on the left, choose **Shared Storage**.
3. On the **Shared Storage** page, select the target cluster from the drop-down list in the upper left corner.
4. In the list, view the file system attached to the current cluster and click **Unmount** in the **Operation** column.
5. Click **Yes**.

3.6 HPC Management and Scheduling Plug-in

3.6.1 Partition Management

Partitions are used to divide cluster nodes into logical resource pools to meet job scheduling requirements (such as by priority, hardware type, and resource limit). There are no limitations on the number of nodes in each partition. Nodes in different partitions are independent of each other.

Partition

- Partition
  - Partitioning is the most basic method for creating resource pools from nodes in a cockpit. Physical nodes are divided into multiple logical resource groups.
  - You can configure resource limits, scheduling policies, and access permissions for each partition.

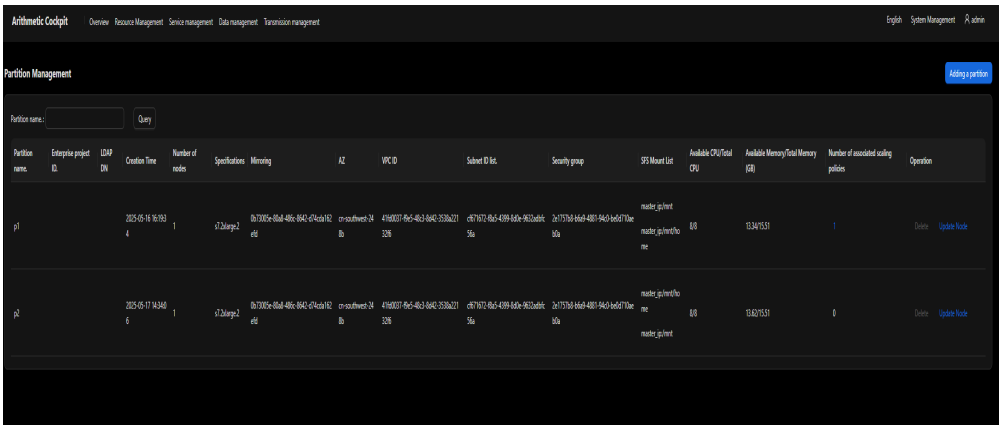

- Node features

  Nodes are tagged by features such as hardware type and network architecture. Logical partitions are further divided based on these features.

 Copyright © Huawei Technologies Co., Ltd.

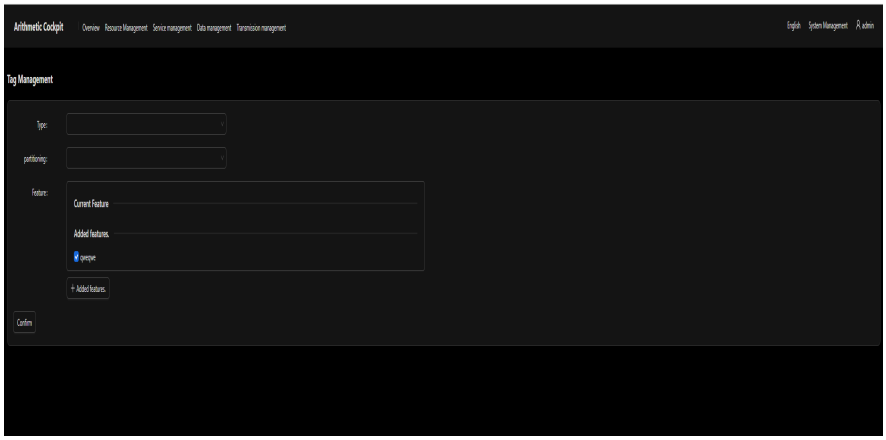

- Dynamic node pools

  Cloud auto scaling or node status management can be used together to dynamically scale a node pool, for example, creating nodes on demand and adding them to a partition.

Configuring a Partition

1. Create a partition.

   Specify parameters such as the image, network, and security group for the partition. Nodes can be dynamically added to or deleted from the partition when the partition is running.

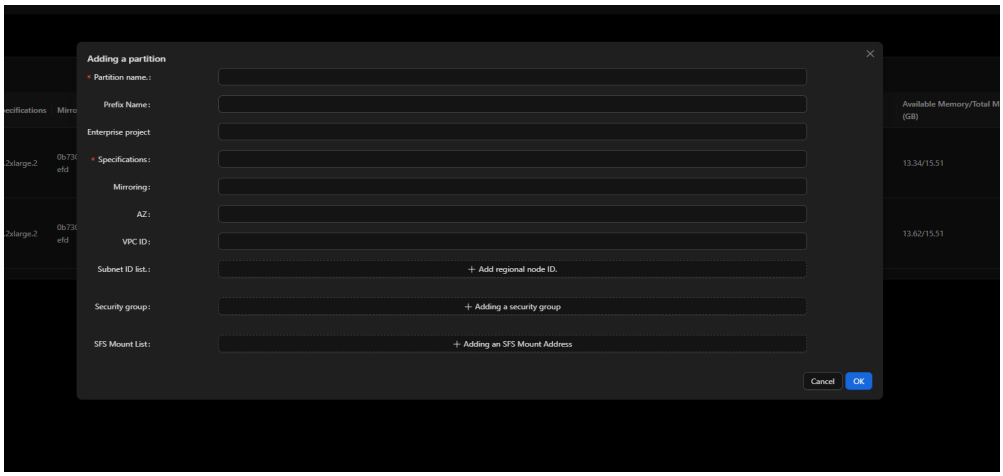

2. Edit the partition.

   Modify the partition of a node.

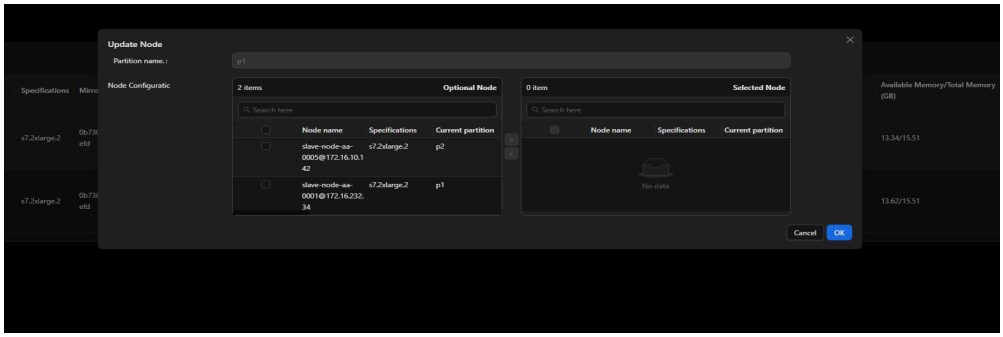

 Copyright © Huawei Technologies Co., Ltd.

3. Submit a job to a partition. For details, see **Job Queues**.

Applying a Scaling Policy to a Partition

Specify a scale-out or scale-in policy on the UI.

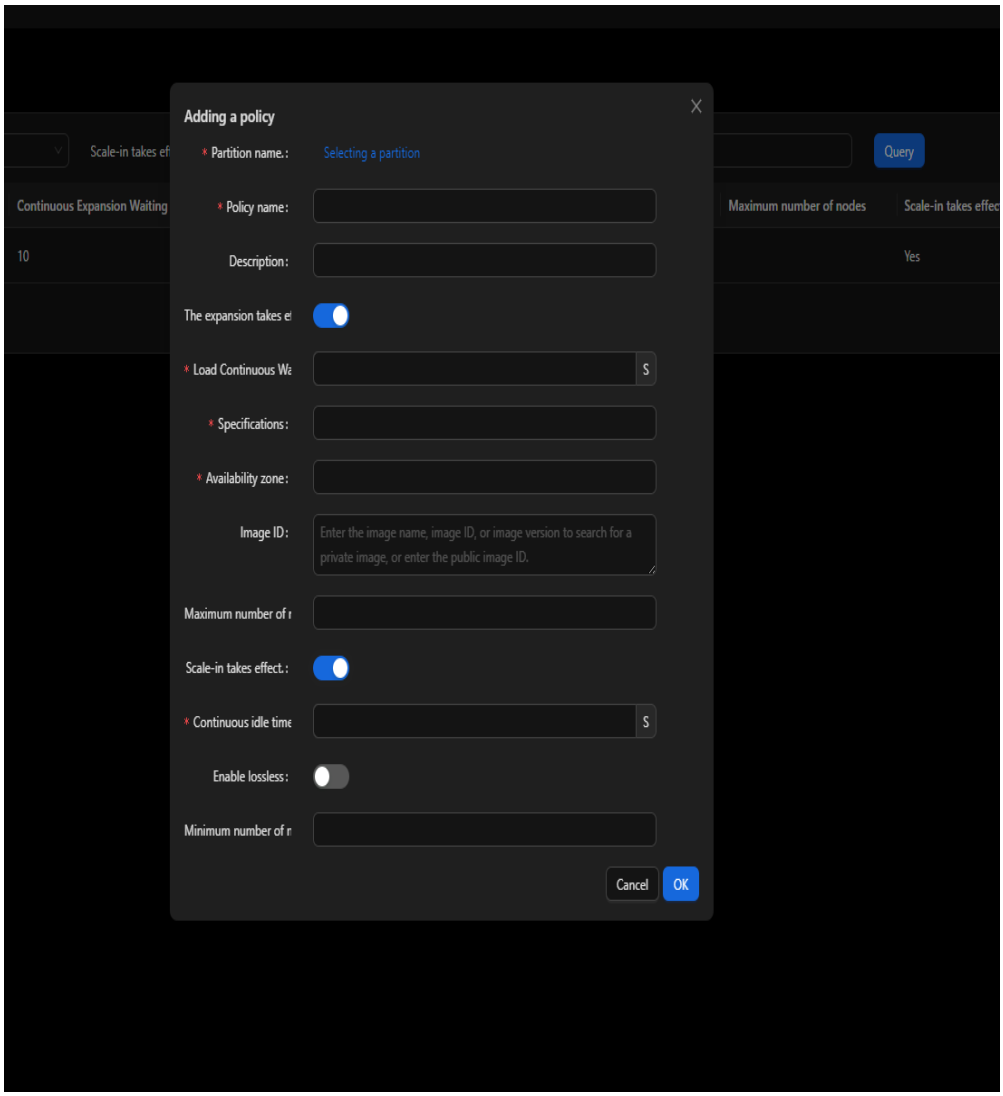

Viewing Partition Details

- List all partitions.

  sinfo -s # Check the partition summary.

  sinfo -p gpu # Check the node statuses in a partition.
- View node details.

  scontrol show nodes node1 # Check the details such as the partition, features, and resources.
- Monitor the partition usage.

  squeue -p gpu -o "%i %P %u %T" # Check the job statuses in the partition named **p1**.
- View partitions.

 Copyright © Huawei Technologies Co., Ltd.

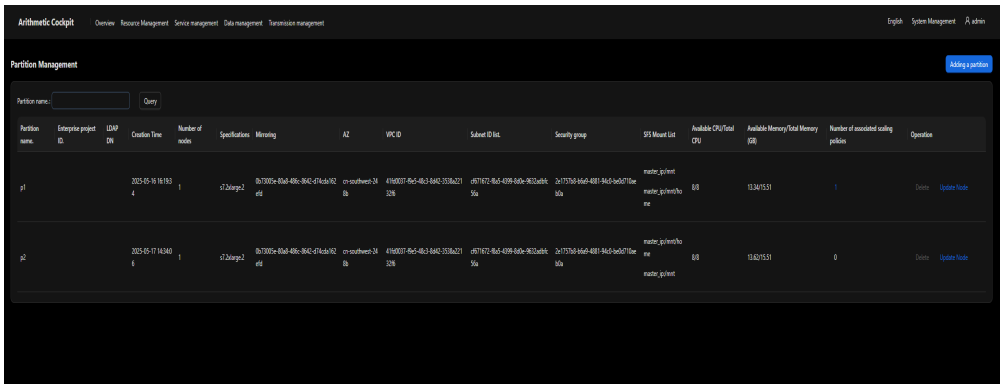

- View the job statuses in a partition.

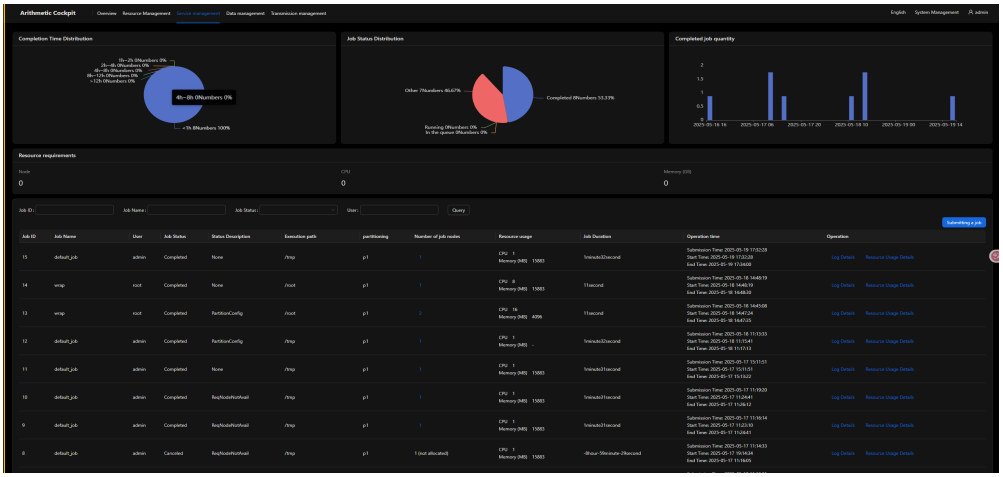

3.6.2 Resource Topology

In Slurm, the **topology.conf** file defines the physical topology of a cluster (such as the rack or switch layer) to optimize job scheduling policies (such as by preferentially allocating adjacent nodes).

The following is a detailed description and example of the configuration file.

Basic Syntax

```
SwitchName=<layer-name>[Nodes=<node-list>][Switches=<sub-layer-list>][Children=<number-of-child-nodes>]
```

- **SwitchName**: defines the name of a topology layer (for example, rack or switch).
- **Nodes**: lists the nodes (for example, node[1-10]) at the topology layer.
- **Switches**: indicates the sub-layers (used for the nested structure) contained in the topology layer.
- **Children**: indicates the number of sub-layers. This field is optional.

Configuration Examples

- Single-layer topology (racks and nodes)

```
# Define two racks (rack1 and rack2). Each rack contains 10 nodes.
SwitchName=rack1 Nodes=node[1-10]
SwitchName=rack2 Nodes=node[11-20]
```

 Copyright © Huawei Technologies Co., Ltd.

- Multi-layer topology (cluster, racks, switches, and nodes)

```
# Cluster (cluster1 is deployed across two equipment rooms.)
SwitchName=cluster1 Switches=room1,room2

# Equipment room (There are two racks in room1.)
SwitchName=room1 Switches=rack1,rack2

# Rack (There are two switches on rack1.)
SwitchName=rack1 Switches=switch1,switch2

# Switch (switch1 connects to 10 nodes.)
SwitchName=switch1 Nodes=node[1-10]
SwitchName=switch2 Nodes=node[11-20]
```

- Hybrid layer (switches and nodes)

```
# The rack houses switches and nodes.
SwitchName=rack1 Nodes=node1,node2 Switches=switch1
SwitchName=switch1 Nodes=node3,node4
```

Key Parameters

Parameter	Description
Nodes	List of continuous or discrete nodes in the format of node[1-5,10].
Switches	List of sub-layer names (used to define the nested structure, such as a rack containing switches).
Children	Number of sub-layers. For example, **Children=2** indicates two sub-layers. This parameter is usually used together with **Switches**.

Verification

1. Check the topology structure.

```
scontrol show topology
```

   Example output:

```
SwitchName=rack1 Nodes=node[1-10]
SwitchName=rack2 Nodes=node[11-20]
```

2. View the topology associated with a node.

```
scontrol show node node01
```

   The **SwitchName** field in the output indicates the topology layer.

```
NodeName=node01 ... SwitchName=rack1
```

3. Submit a test job.

```
# Allocate two nodes on the same rack.
sbatch --nodes=2 --switches=1@rack1 --wrap="hostname"
squeue -o "%N" # Check whether the allocated nodes are on rack1.
```

FAQ

- A node is not associated with any layer.
  - **Symptom**: **SwitchName** in the **scontrol show node** command output is empty.
  - **Solution**: Check whether the **Nodes** parameter contains the node.

 Copyright © Huawei Technologies Co., Ltd.

- There are syntax errors.
  - **Symptom**: slurmctld fails with an error message "invalid topology.conf" displayed in the log.
  - **Solution**: Run the **scontrol reconfigure** command to reload the configuration and then view the log.
    ```
    tail -f /var/log/slurm/slurmctld.log
    ```
- Failed to allocate nodes for a job.
  - **Reason**: The value of the **--switches** parameter exceeds the capacity of the actual layer.
  - **Example**: If rack1 has only 10 nodes, submitting **--nodes=12 --switches=1@rack1** will fail.

Best Practices

- Align with the actual hardware.

  Ensure that the hierarchy (such as racks and switches) in **topology.conf** is consistent with the physical environment configuration.
- Reduce the number of layers.

  Do not nest too many layers (for example, clusters, equipment rooms, racks, switches, and nodes). Generally, two or three layers are enough.
- Limit the use of topology resources based on QoS.

  Use QoS to limit the resources that can be used by users at a specific topology layer.
  ```
  sacctmgr add qos high_priority --max-switches-per-job=2
  ```

Configuration Example (Dynamic Resource Partitioning)

```
# Define the independent topology of GPU nodes.
SwitchName=gpu_rack Nodes=gpu[1-4] # A dedicated rack for GPU nodes
# Specify the GPU topology when submitting a job.
sbatch --gres=gpu:2 --switches=1@gpu_rack job.sh
```

3.6.3 Job Queues

A job queue is a collection of jobs that are running or waiting for resource allocation. Each partition corresponds to a logical queue. You can set independent scheduling policies and resource limits for each partition.

Partition

For details about partitions, see **Partition Management**.

Managing Job Queues

- Set the CPU and memory of a job and submit the job to a specific partition.

 Copyright © Huawei Technologies Co., Ltd.

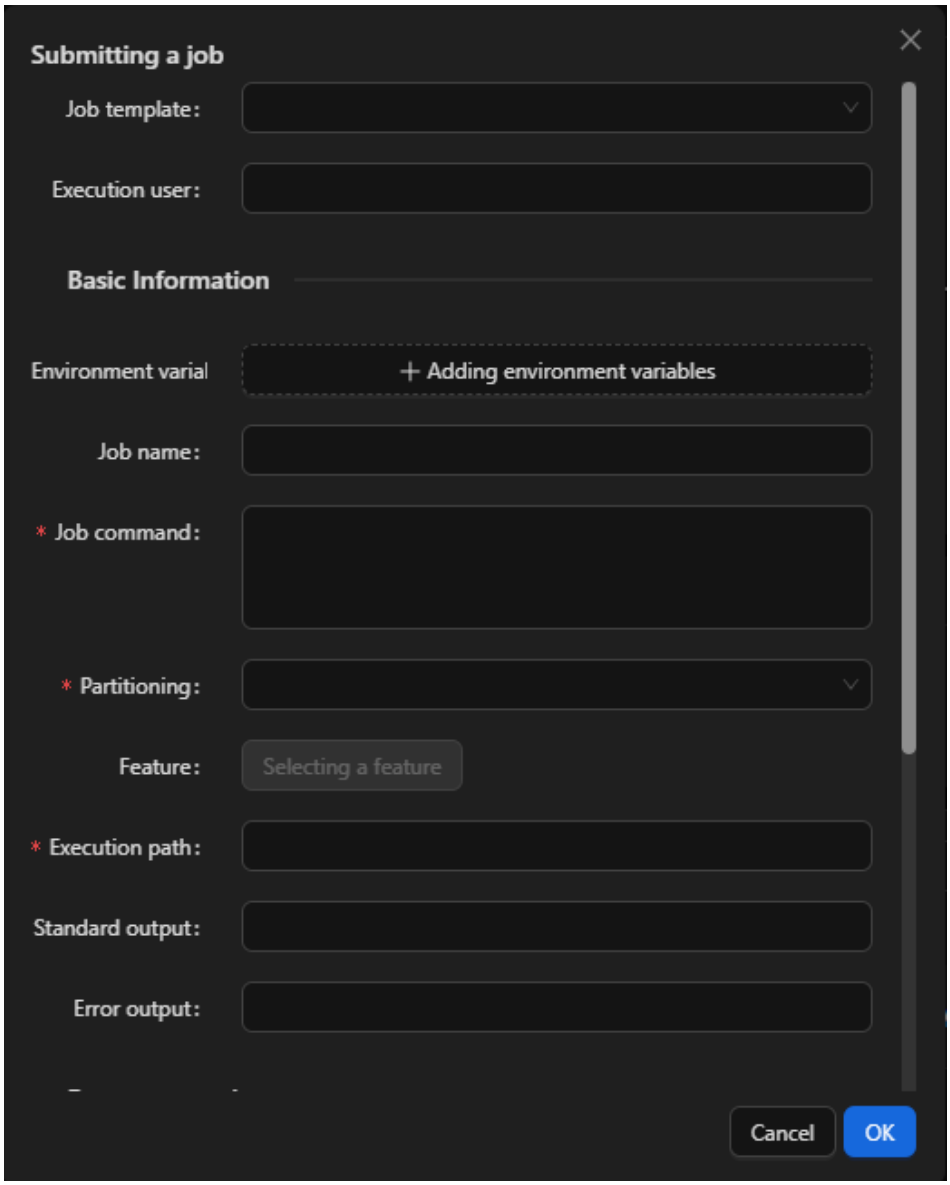

- View the job statuses.

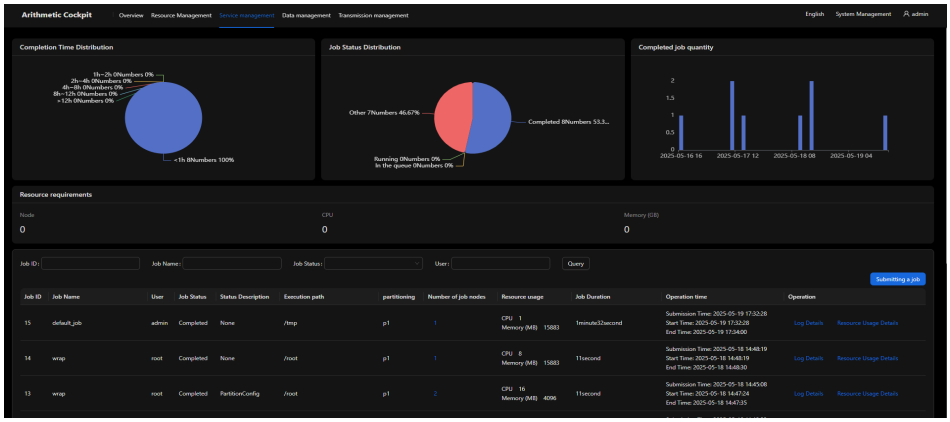

Job Templates

Create a job template to set the CPU and memory for jobs.

 Copyright © Huawei Technologies Co., Ltd.

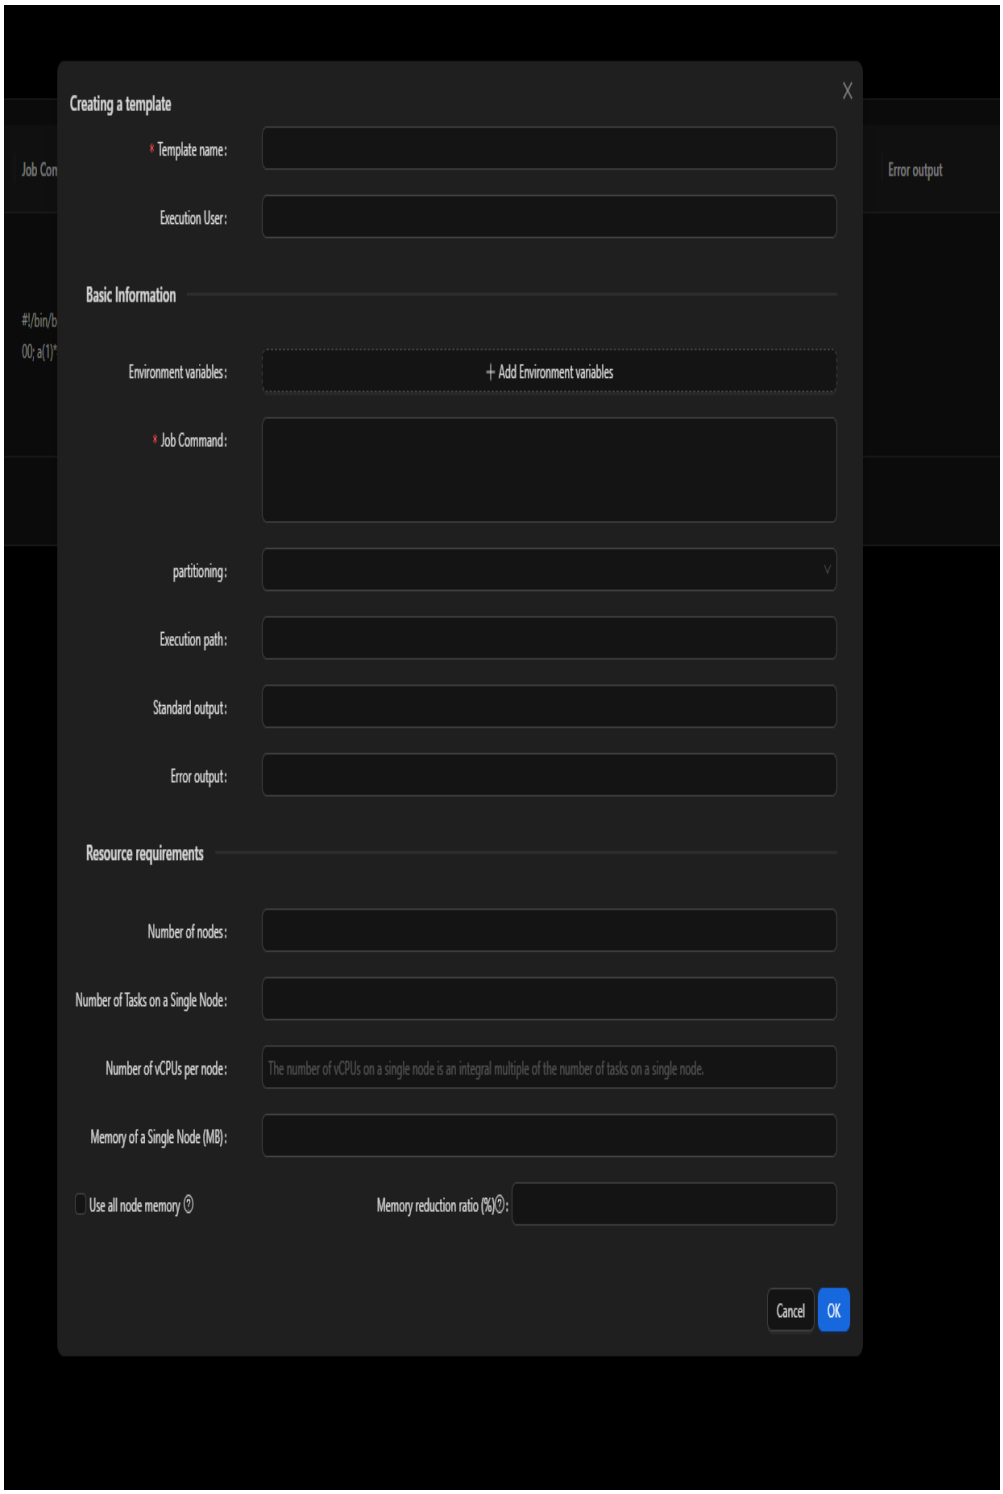

QoS

QoS is used to limit the resource usages (such as the number of concurrent jobs and maximum CPU and memory) for users or groups. It is used together with partitions for fine-grained control.

1. Create a QoS policy.
   ```
   sacctmgr add qoshigh_priority set MaxJobs=50 Priority=1000
   ```

 Copyright © Huawei Technologies Co., Ltd.

- **MaxJobs**: indicates the maximum number of jobs.
- **Priority**: A larger value indicates a higher priority.

```
[root@master ~]# sacctmgr add qos high1 set MaxJobs=100 Priority=1000
 Adding QOS(s)
  high1
 Settings
  Description    = high1
  MaxJobsPerUser = 100
  Priority               = 1000
Would you like to commit changes? (You have 30 seconds to decide)
(N/y): y
[root@master ~]#
[root@master ~]#
```

2. Apply the QoS policy to a partition.

```
PartitionName=p1 Nodes=slave-node-aa-0000 State=UP AllowQos=high1
PartitionName=p2 State=UP
```

Jobs can only use the QoS set for the partition.

Job Functions

- Job Dependencies

sbatch --dependency=afterok:<job_id> job2.sh # Run job2 after job1 is successfully executed.

```
[root@master sfs_turbo_1]# sbatch --dependency=after:118_0 cpu_sbatch.sh
Submitted batch job 123
[root@master sfs_turbo_1]# squeue
             JOBID PARTITION     NAME     USER ST       TIME  NODES NODELIST(REASON)
             118_4       x86       ty     root PD       0:00      1 (Resources)
               123       x86       ty     root PD       0:00      1 (None)
             118_3       x86       ty     root  R       0:56      1 hpc-aa-0009
             118_0       x86       ty     root  R       1:35      1 hpc-aa-0009
             118_1       x86       ty     root  R       1:35      1 hpc-aa-0008
             118_2       x86       ty     root  R       1:35      1 hpc-aa-0008
```

- Job Array

sbatch --array=0-4 job_array.sh # Submit 100 similar jobs.

```
[root@master sfs_turbo_1]# sbatch --array=0-4 cpu_sbatch.sh
Submitted batch job 118
[root@master sfs_turbo_1]# squeue
             JOBID PARTITION     NAME     USER ST       TIME  NODES NODELIST(REASON)
         118_[3-4]       x86       ty     root PD       0:00      1 (Resources)
             118_0       x86       ty     root  R       0:01      1 hpc-aa-0009
             118_1       x86       ty     root  R       0:01      1 hpc-aa-0008
             118_2       x86       ty     root  R       0:01      1 hpc-aa-0008
               117       x86 default    admin  R       1:11      1 hpc-aa-0009
```

3.6.4 Quota Management

Quota management is used to limit the use of compute resources (such as CPUs, memory, nodes, and jobs) by users and partitions. Slurm does not directly manage storage or disk quotas. However, it can control compute resource quotas through Quality of Service (QoS), associations, and users/partitions.

Core Concepts of Quota Management

1. Backend configuration
   - Association
     - Association is the core mechanism that defines the relationships between users and QoS in Slurm.

 Copyright © Huawei Technologies Co., Ltd.

- Using association, you can set resource limits (such as the maximum number of jobs and maximum number of CPUs or nodes) for specific users.
- The table below describes the configuration parameters.

Parameter	Description
MaxJobs	Maximum number of jobs allowed to run
MaxCPUs	Maximum number of CPU cores
MaxNodes	Maximum number of nodes
MaxSubmit	Maximum number of jobs a user can submit
MaxWall	Maximum wall clock time a job can run (by partition or QoS)

– Quality of service (QoS)

- QoS is used to define resource limits and priorities for jobs. It can be associated with users, accounts, or jobs.
- You can set the following parameters using QoS.

Parameter	Description
Priority	Determines the scheduling order of jobs.
MaxWall	Sets the maximum wall clock time a job can run.
MaxJobs/MaxCPUs	Defines resource limits (the maximum number of jobs or the maximum number of vCPUs).

2. Settings on the UI

   You can set user quotas on the UI.

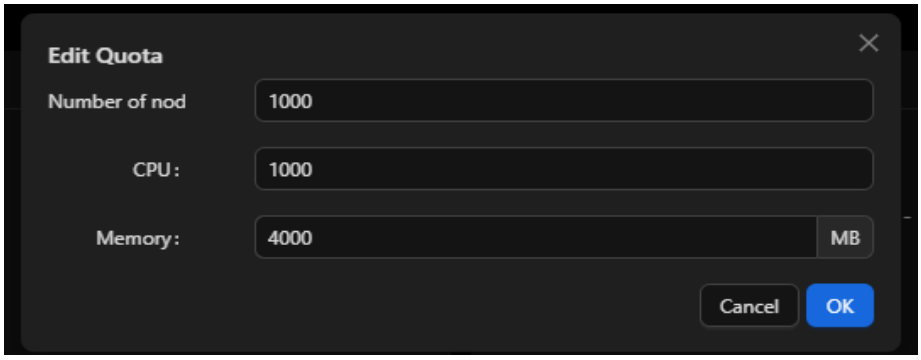

 Copyright © Huawei Technologies Co., Ltd.

Quota Types and Configuration Methods

1. Walltime limit

   The walltime limit defines the maximum amount of time a job can run before it is automatically terminated. You can enforce the walltime limit at multiple levels to control how long a job can run:

   – QoS level: Set **MaxWall** for QoS by running the **sacctmgr** command.

   # Example: Create a QoS with the walltime limit set to 24 hours.

   ```
   sacctmgr modify qos normal set MaxWall=24:00:00
   ```

   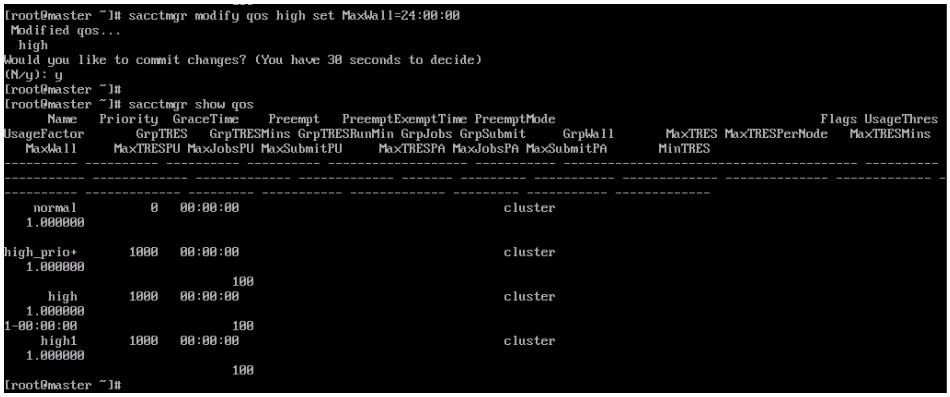

   - User level:

   ```
   sacctmgr modify user alice set MaxWall=2-00:00:00
   ```

2. Resource quotas (CPUs/Nodes/Jobs)

   – QoS level:

   ```
   acctmgr modify qos normal set MaxCPUs =100 MaxJobs=100
   ```

   – User level:

   Use sacctmgr to configure various resource limits, including MaxJobs, MaxCPUs, and MaxNodes.

   ```
   # Set the maximum number of jobs to 10 and the maximum number of CPUs to 100 for user alice.
   sacctmgr modify user alice set MaxJobs=10 MaxCPUs=100
   ```

   Settings on the UI

   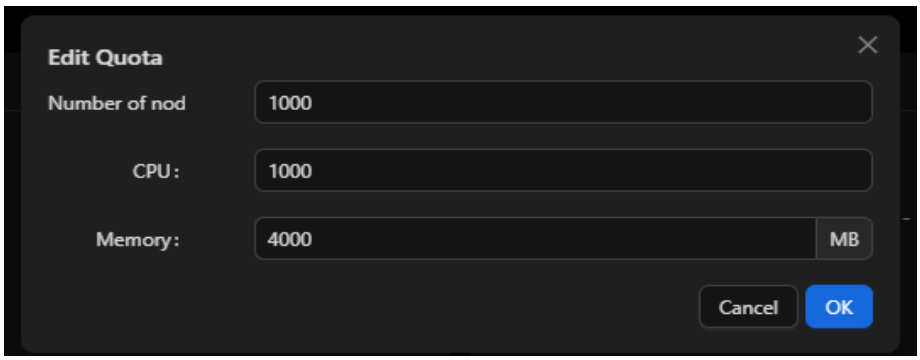

Configuration Examples

Example 1: Setting QoS and resource limits for a user

1. Create a QoS.

   ```
   sacctmgr add qos name=short_qos MaxWall=1:00:00 MaxJobs=5
   ```

 Copyright © Huawei Technologies Co., Ltd.

2. Associate the QoS with the user.
   sacctmgr modify user alice set qos=short_qos
3. Verify the settings.
   sacctmgr show user alice format=User,Account,QOS,MaxJobs,MaxCPUs

```
[root@master ~]# sacctmgr add qos qos_admin set Priority=100 MaxWall=7-00:00:00 MaxJobs=50 MaxNodes=10
 Adding QOS(s)
  qos_admin
 Settings
  Description    = qos_admin
  MaxJobsPerUser = 50
  MaxTRESPerJob           = node=10
  MaxWall                 = 7-00:00:00
  Priority                = 100
Would you like to commit changes? (You have 30 seconds to decide)
(N/y): y
[root@master ~]#
[root@master ~]# sacctmgr modify user admin set qos=qos_admin
 Modified user associations...
  C = intelligentcockpit A = root                 U = admin
Would you like to commit changes? (You have 30 seconds to decide)
(N/y): y
[root@master ~]#
[root@master ~]# sacctmgr list associations user=admin format=User,GrpTRES -p -n
admin|cpu=4,mem=4000M,node=10|
```

Example 2: Denying requests that exceed the quota

```
[root@master ~]# squeue
             JOBID PARTITION     NAME     USER ST       TIME  NODES NODELIST(REASON)
                28        p1 default_    admin PD       0:00      1 (AssocGrpMemLimit)
```

Checking the Quota

- Check associations.
  sacctmgr show assoc # Displays all associations.
  sacctmgr show user alice # Displays the restrictions for user **alice**.
- Check QoS settings.
  sacctmgr show qos
- Check the resource usages.
  sshare -a #Display the resource usages of the account.
- Resource usages on the UI

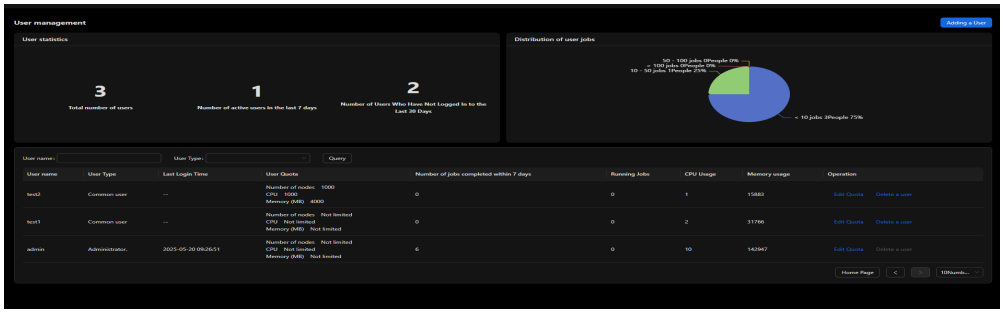

Precautions

1. Priority: In Slurm, user-level quotas generally take precedence over QoS-level quotas.
2. Audit and monitoring: You can use **sshare** and **sacct** to monitor resource usages periodically.

Remarks

You can view the quota usage in the last 14 days.

 Copyright © Huawei Technologies Co., Ltd.

3.6.5 Dynamic Scheduling

Dynamic scheduling optimizes job execution by adapting to real-time cluster conditions, ensuring efficient resource utilization while maintaining fair access for all users.

Dynamic Resource Scheduling

Function 1: Scaling up when loads stack up and scaling down when nodes are idle

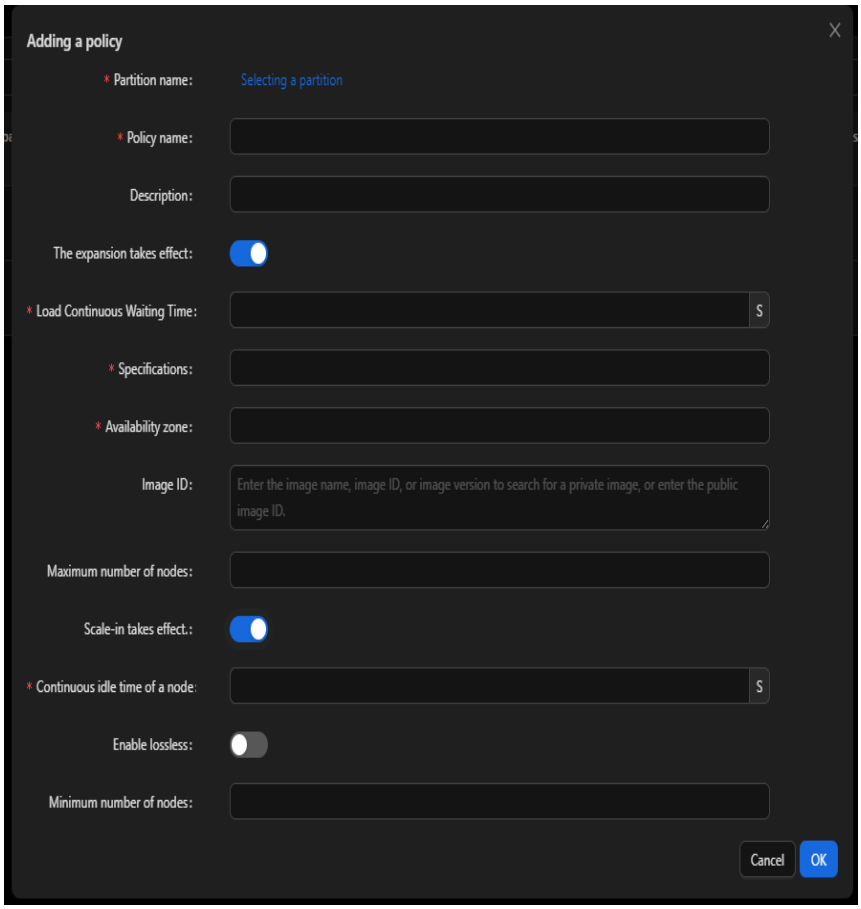

- If no node is available, new sites are automatically added at the preset time.
- Resources are automatically released when a node remains idle for the preset time (300 seconds).

Function 2: Setting stable nodes

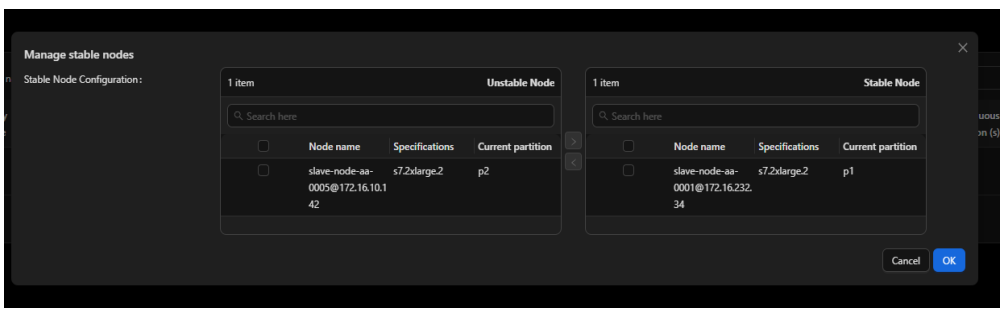

 Copyright © Huawei Technologies Co., Ltd.

No nodes are deleted even when they are idle.

Function 3: Provisioning hybrid resources elastically

You can create compute nodes of different specifications in a partition.

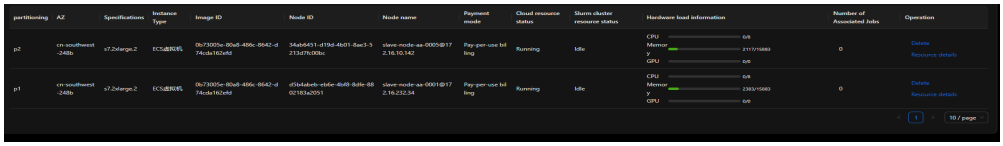

Checking resource scheduling

Check the partitions where the compute nodes are located.

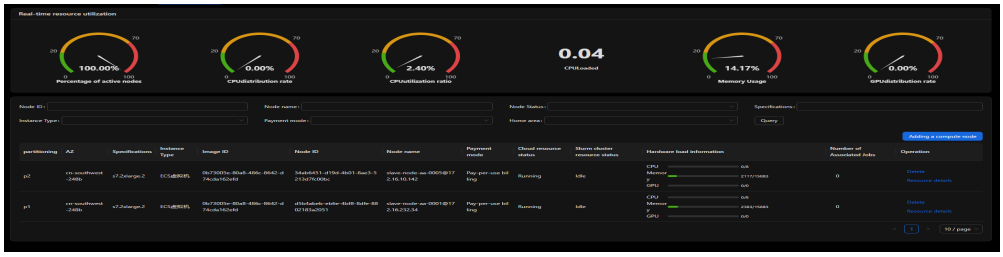

Dynamic Job Scheduling

Job Scheduling Policies

1. Backfill scheduling
   - Principles: Low-priority small jobs are executed when resources are idle without delaying high-priority jobs.
   - Dynamic performance of the HPC Management and Scheduling Plug-in:
     - Calculates the earliest start time (EST) of each job in real time.
     - Scans queues to schedule jobs for idle resources to process.
   - Key dependency: You need to specify the accurate start time (in the **--time** parameter). Otherwise, the backfilling efficiency decreases.
   - Example:
     ```
     sbatch -p [partition] --time=00:30:00 job.sh # Accurately specify the job start time to facilitate backfilling.
     ```
2. Preemption
   - Principles: When a high-priority job is submitted, Slurm can terminate or suspend low-priority jobs to release resources for use by this high-priority job.
   - Preemption mechanisms

Type	Action	Application Scenarios
Requeue	When a job is preempted, it is returned to the queue to be scheduled again later.	Long jobs that can be interrupted

 Copyright © Huawei Technologies Co., Ltd.

Type	Action	Application Scenarios
Suspend	A job is paused, allowing higher-priority jobs to run first. Once resources become available, the suspended job resumes.	Running of short-term high-priority jobs
Cancel	The job is terminated completely and does not get rescheduled.	Emergency tasks (exercise caution when setting this value)

- **Configuration methods**:

  Enable the preemption policy in the QoS or partition.

```
# Create a QoS preemption.
sacctmgr add qos low_prio Preempt=high_prio # Allow jobs running under the low_prio QoS to be preempted by jobs associated with the high_prio QoS.
```

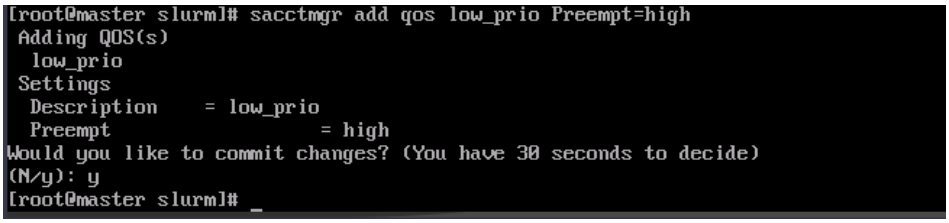

3. Dynamic adjustment of the priority
   - Fair-share: The priority weights are dynamically calculated based on historical resource usages of users or user groups (for example, CPU hours used in the last 30 days).

     sshare -l # Check the fair-share values of users. (A smaller fair-share value means a higher scheduling priority.)
   - Real-time feedback mechanism:
     - When a user uses more resources than allowed, the user's job priority automatically decreases. Otherwise, the user's job priority increases.
     - The weight is calculated using the following formula: Priority = f(Age, FairShare, QoS, Partition,...). The priority is recalculated every several seconds.

Dynamical scheduling optimization policies

1. On the administrator side
   - **Enable backfill scheduling**:

     # Enable backfill scheduling in the **slurm.conf** file.

```
SchedulerType=sched/backfill # Use the backfill scheduler.
SchedulerParameters=bf_max_job_test=100 # Set the maximum number of jobs that can be scanned during each scheduling.
```

   - **Configure the preemption policy**.

```
# Define global preemption rules in slurm.conf.
PreemptType=preempt/partition_prio # Preempt by partition priority.
PreemptMode=SUSPEND                # Suspend jobs, instead of terminating them.
```

 Copyright © Huawei Technologies Co., Ltd.

- **Set dynamic priority weights**.
  ```
  PriorityType=priority/multifactor   # Enable multi-factor priority calculation.
  PriorityWeightFairshare=1000        # Specify the weight of fair-share usage.
  PriorityWeightAge=100               # Specify the weight of the job age.
  ```

2. On the user side
   - **Accurately specify resource requirements**.
     ```
     sbatch -p [partition] --ntasks=4 --mem=8G --time=1:00:00 job.sh # Avoid requesting excessive resources.
     ```
   - **Use QoS to improve priority**.
     ```
     sbatch --qos=urgent job.sh # Submit the job to a high-priority QoS queue.
     ```

Job scheduling monitoring

1. Check the backfill scheduling.
   ```
   scontrol show config | grep Backfill # Check the backfill scheduler status.
   sdiag -f | grep "Backfill cycle"     # Check the backfill scheduling cycles.
   ```
2. Analyze preemption events.
   ```
   sacct -X -j <job_id> -o JobID,PreemptedBy # Check the jobs whose resources are preempted by other jobs.
   ```
3. Track the priority in real time.
   ```
   sprio -w -j <job_id> # Displays a detailed breakdown of the priority factors influencing a specific job's scheduling.
   ```
4. 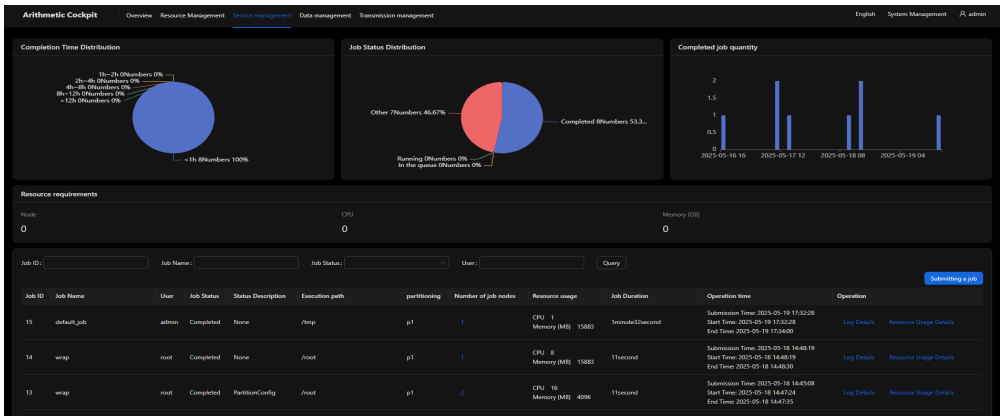

Typical Scenarios and Solutions

Scenario 1: Small jobs are queued for a long time.

- **Problem**: Small jobs experience delays when large jobs occupy resources.
- **Solution**: Enable backfill scheduling with the **--time** parameter correctly specified.
  ```
  sbatch --time=00:10:00 small_job.sh # Specify the running time to help backfill.
  ```

Scenario 2: An urgent job needs to be executed immediately.

- **Problem**: High-priority jobs need to preempt resources for running immediately.
- **Solution**: Add a QoS with preemption enabled.
  ```
  # Administrator configuration
  sacctmgr add qos urgent Priority=1000 Preempt=low_prio
  # Submitted by users
  sbatch --qos=urgent critical_job.sh
  ```

Scenario 3: Uneven resource usages

 Copyright © Huawei Technologies Co., Ltd.

- **Problem**: Some nodes are idle, but jobs are still queuing.
- **Solution**: Check whether the partition configuration is too strict or adjust the node sharing policy (for example, enable **OverSubscribe**).

Limitations of Dynamic Scheduling

- **Time estimation deviation**: If the user underestimates the job running time, backfilling will fail. If the user overestimates the job running time, resources will be wasted.
- **Preemption cost**: Frequent job suspension or restart may increase the system overhead (for example, when the checkpoint mechanism is disabled).
- **Configuration complexity**: Dynamic scheduling parameters need to be finely adjusted (such as weight allocation and backfilling depth). Improper configuration may reduce the overall efficiency.

Summary

You can set automatic scaling policies to schedule resources to maximize resource utilization. Dynamic job scheduling implements efficient resource utilization through backfilling, preemption, and real-time priority adjustment. You should properly specify the resource requirements and time limits. The administrator needs to optimize scheduling policy parameters (such as the weights and preemption rules).

3.6.6 Capacity Reservation

Resource reservation allows users or administrators to reserve compute nodes to ensure that resources can be exclusively used or shared in a specified period. The following describes how to reserve resources.

Creating Resource Reservations

Run the **scontrol create reservation** command (as the administrator):

```
scontrol create reservation ReservationName=<name> StartTime=<start-time> Duration=<duration>
Nodes=<node-list-or-quantity> Partition=<partition-name> Users=<user-list> Flags=<options>
```

Parameter description:

- **ReservationName**: specifies the unique name of the reservation record.
- **StartTime**: defines when the reservation begins. The format is YYYY-MM-DD[THH:MM[:SS]] or now+{number}{unit} (for example, now+1hours).
- **Duration**: specifies how long the reservation lasts. The format is DD-HH:MM:SS (for example, 2-12:00:00 indicates 2 days and 12 hours).
- **Nodes**: determines which nodes are included.
- **Partition**: determines which partition the reservation applies to.
- **Users**: assigns the reservation to specific users. (Multiple users are separated by commas, for example, user1,user2. The creator can access the reservation by default).
- **Flags**: controls how resources are allocated.
  - MAINT: marks the reservation for maintenance purposes. Only administrators can access it.

 Copyright © Huawei Technologies Co., Ltd.

  - OVERLAP: allows the reservation to overlap with other reservations.
  - IGNORE_JOBS: allows the reservation to be created without considering currently running jobs and to take effect immediately.
  - SPEC_NODES: ensures that specific nodes are reserved, rather than allowing Slurm to allocate any available nodes.

Create a capacity reservation policy in the backend of the cockpit (to reserve a compute node for user test1).

```
[root@master ~]# scontrol create reservation ReservationName=huawei_reserve StartTime=2025-05-17T10:59:00 Duration=1:00:00 Nodes
=slave-node-aa-0000 Users=test1
Reservation created: huawei_reserve
[root@master ~]#
```

Checking Resource Reservation

- Check all reservations.
  scontrol show reservation
- Check a specific reservation:
  scontrol show reservation <ReservationName>

Check the capacity reservation policy in the backend of the cockpit.

```
[root@master ~]# scontrol show reservation
ReservationName=huawei_reserve StartTime=2025-05-17T10:59:00 EndTime=2025-05-17T11:59:00 Duration=01:00:00
   Nodes=slave-node-aa-0000 NodeCnt=1 CoreCnt=4 Features=(null) PartitionName=(null) Flags=SPEC_NODES
   TRES=cpu=8
   Users=test1 Groups=(null) Accounts=(null) Licenses=(null) State=ACTIVE BurstBuffer=(null) Watts=n/a
   MaxStartDelay=(null)

[root@master ~]#
```

Managing Reservations

- Update reservation parameters (for example, modifying the duration).
  scontrol update ReservationName=<*Name*> Duration=<*New duration*>

  ```
  [root@master ~]# scontrol update ReservationName=huawei_reserve Duration=00:40:00
  Reservation updated.
  [root@master ~]# _
  ```

- Delete a reservation.
  scontrol delete ReservationName=<*name*>

  ```
  [root@master ~]# scontrol delete ReservationName=huawei_reserve
  [root@master ~]# scontrol show reservation
  No reservations in the system
  [root@master ~]# _
  ```

Preventing Non-specified Users from Using Reserved Resources

You can submit a job in either of the following ways (it is recommended that you use the cockpit):

- **Submitting a job on the cockpit UI**

 Copyright © Huawei Technologies Co., Ltd.

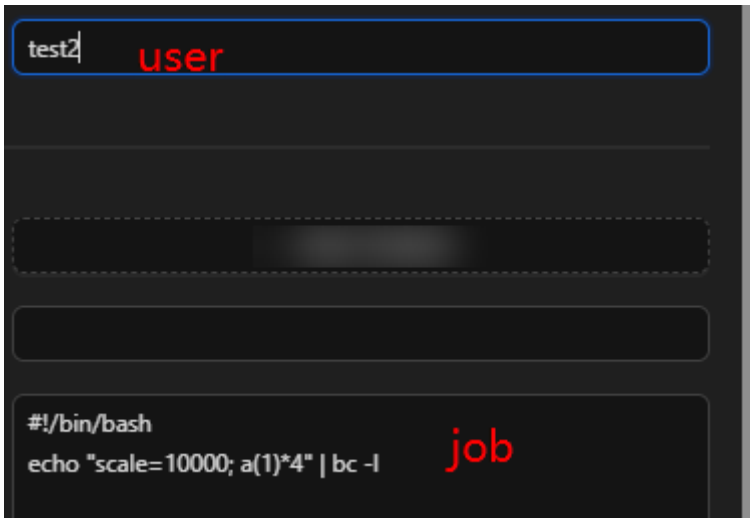

 Copyright © Huawei Technologies Co., Ltd.

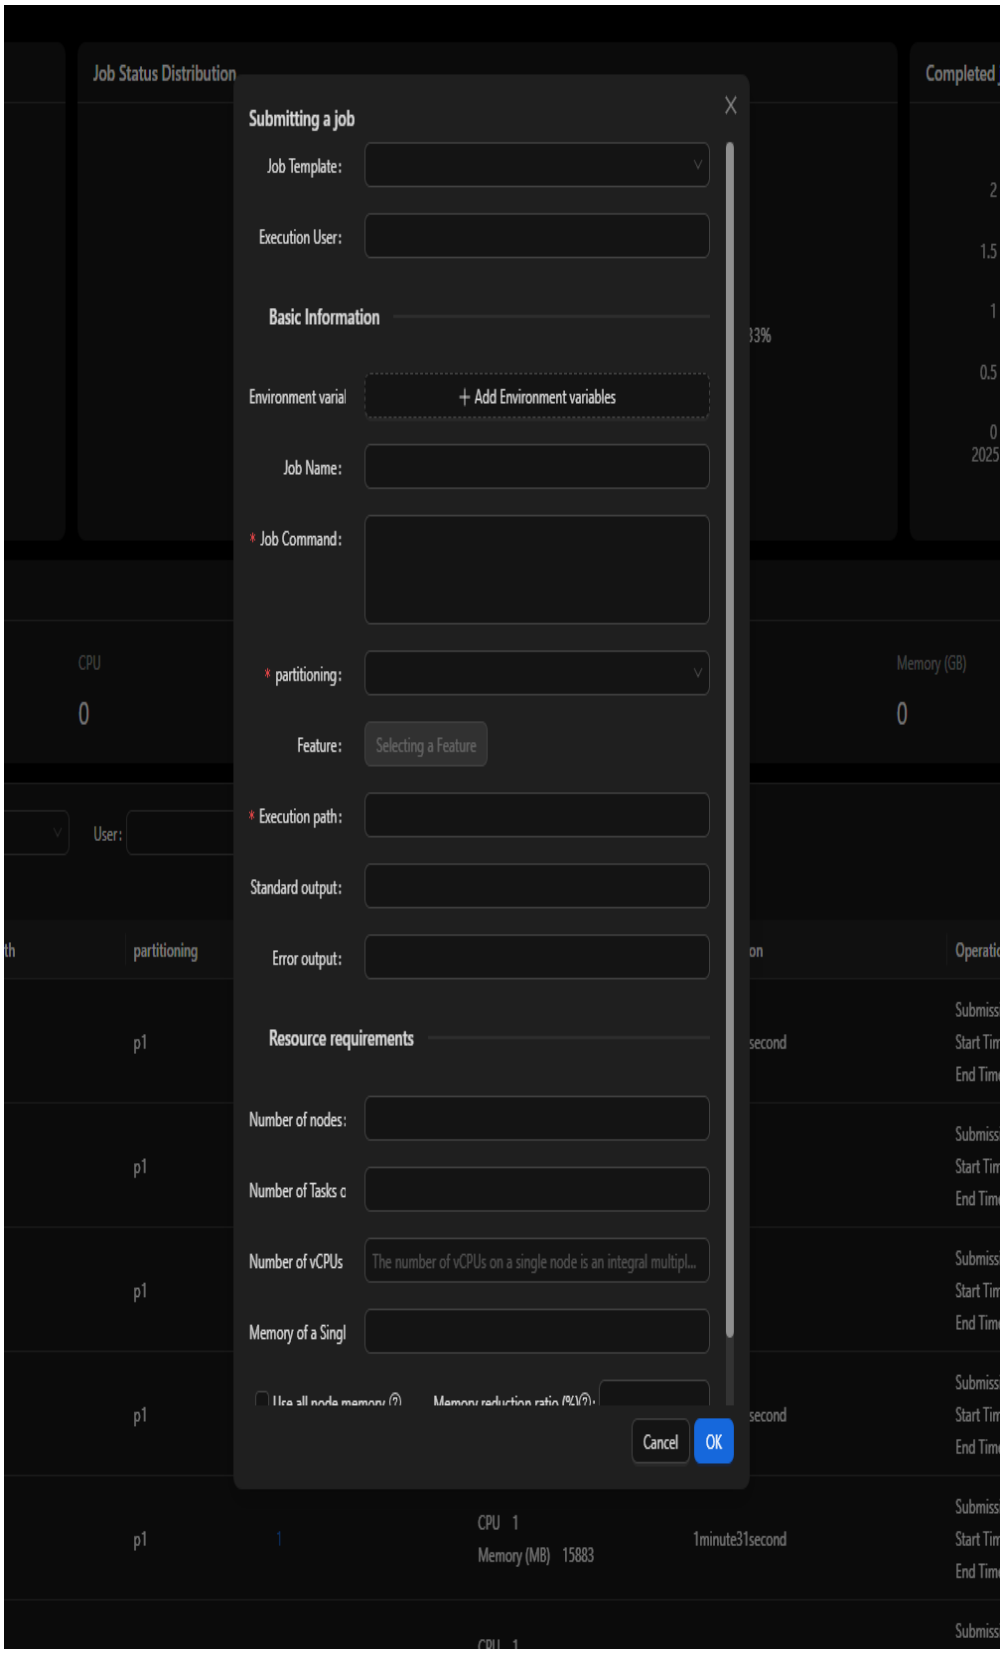

If the user who submits the job is not the specified user, the job cannot be executed.

```
[root@master ~]# squeue
             JOBID PARTITION     NAME     USER ST       TIME  NODES NODELIST(REASON)
                 6        p1 default_    test2 PD       0:00      1 (ReqNodeNotAvail, May be reserved for other job)
[root@master ~]#
```

 Copyright © Huawei Technologies Co., Ltd.

- **Submitting a job by running a command**

  Use **--reservation** in the job submission command to specify the reservation name.

  ```
  # Submit batch jobs.
  sbatch --reservation=<name> --partition=<partition> job.sh
  ```

Use Cases

Scenario 1: Reserve a node for testing.

```
scontrol create reservation \
    ReservationName=test_job \
    StartTime=now+30minutes \
    Duration=1:00:00 \
    Nodes=node01 \
    Users=alice
```

- User **alice** can exclusively use node01 for 1 hour 30 minutes later.

Scenario 2: Reserve nodes for multi-user collaborative tasks.

```
scontrol create reservation \
    ReservationName=team_project \
    StartTime=2024-01-01T09:00:00 \
    Duration=24:00:00 \
    Nodes=4 \
    Partition=workq \
    Users=user1,user2,user3
```

- Users share four nodes for 24 hours in a specified period.

Scenario 3: Reserve nodes for maintenance purposes.

```
scontrol create reservation \
    ReservationName=maintenance \
    StartTime=now \
    Duration=8:00:00 \
    Nodes=ALL \
    Flags=MAINT,IGNORE_JOBS
```

- The administrator occupies all nodes for maintenance and terminates the existing jobs.

Precautions

- **Permissions**: Common users need to create reservations with the help of administrators.
- **Time conflicts**: By default, nodes cannot be reserved repeatedly unless **OVERLAP** is specified.
- **Job time limits**: Jobs must start within the reserved time window and cannot run beyond the reserved duration.

Resource reservation can optimize cluster resource allocation and ensure that critical tasks are executed on time. You are advised to contact the cluster administrator to configure resource reservation when there are complex requirements.

 Copyright © Huawei Technologies Co., Ltd.

3.6.7 bursting from on-premises Plug-in

Introduction

The bursting from on-premises plug-in can communicate with Slurm and Huawei Cloud Auto Scaling (AS), Cloud Eye, and Elastic Cloud Server (ECS) services. This plug-in dynamically adjusts the number of Slurm nodes as needed.

Preparations

1. Creating an AS configuration

   Log in to the AS console, and create an AS configuration based on the required ECS configurations.

   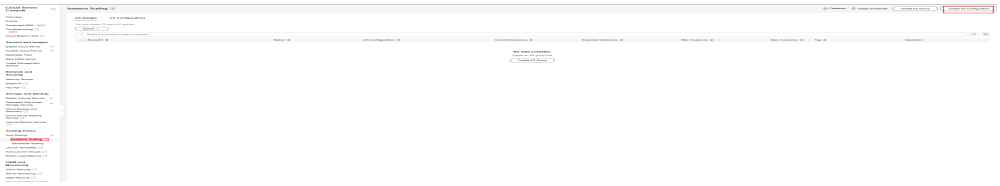

2. Creating an AS group

   Create an AS group using the AS configuration you created.

   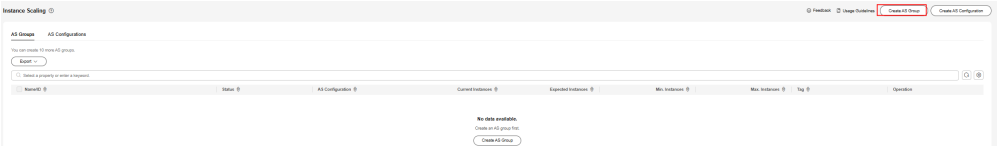

   Record the AS group ID, which will be used in the Scaler configuration file.

   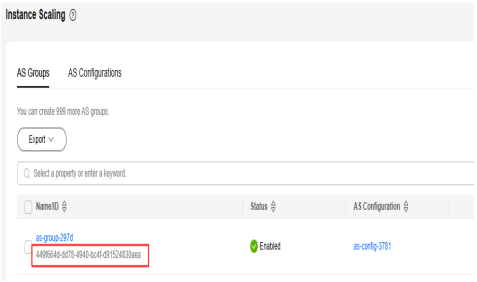

3. Purchasing a yearly/monthly ECS

   Log in to the ECS console, purchase a yearly/monthly ECS as needed, and record the ECS node name in the Slurm cluster. This node will be used as a stable node in the Scaler configuration file.

4. Creating an access key pair

   a. On the management console, choose **My Credentials**.

 Copyright © Huawei Technologies Co., Ltd.

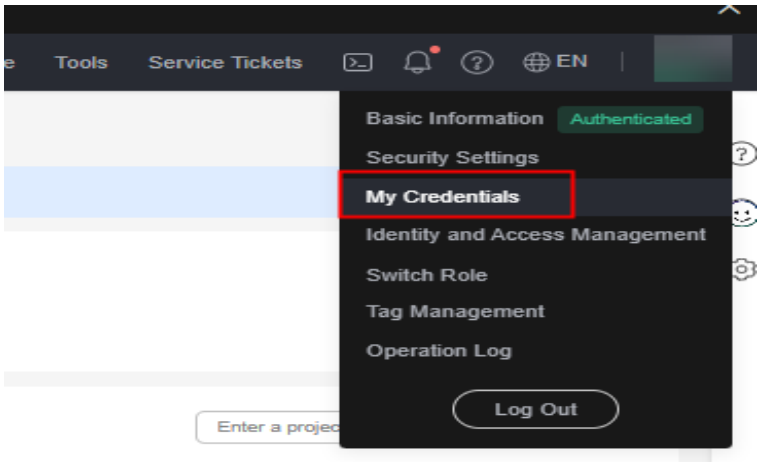

b. In the navigation pane, choose **Access Keys** and click **Create Access Key**. After the access keys are created, you can download the file that contains the created Access Key ID (AK) and Secret Access Key (SK).

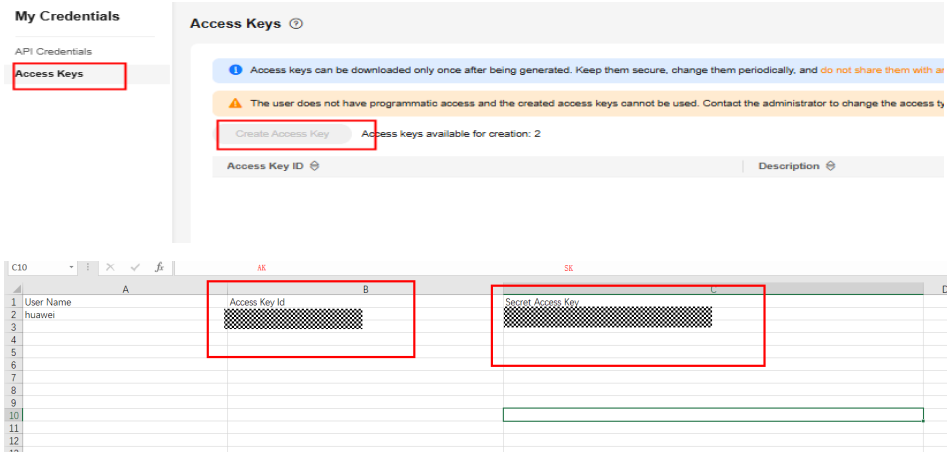

5. Determining the region ID and the project ID

   Go to **My Credentials** page, choose **API Credentials** in the navigation pane, and find the target region ID and the project ID.

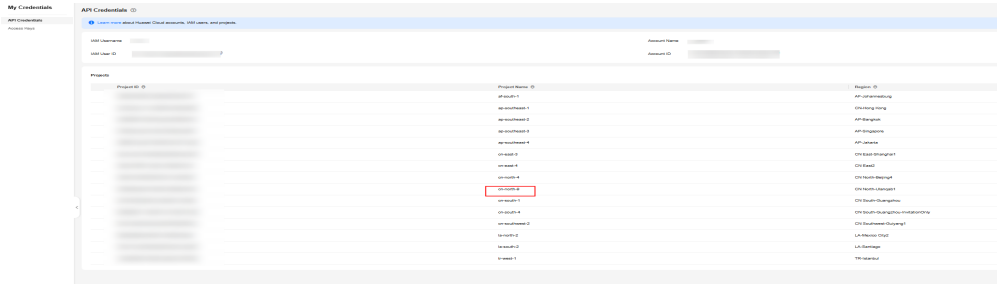

6. Determining the endpoints of antecedent services

   The Scaler program depends on AS, Cloud Eye, and ECS. The endpoints of each service are as follows:

   as.*{Region ID}*.myhuaweicloud.com

   ces.*{Region ID}*.myhuaweicloud.com

   ecs.*{Region ID}*.myhuaweicloud.com

   Replace *{Region ID}* with the actual region ID, for example, as.*cn-north-9*.myhuaweicloud.com.

 Copyright © Huawei Technologies Co., Ltd.

Constraints

- Slurm user quotas can only be configured at the node level.
- Instances cannot be manually added to an AS group.
- To update the stable node information, you need to update the Scaler configuration file, restart the Scaler, and then create an instance.
- A shortened stable node name must be in the format of a prefix followed by a number, like **hpc-0001** and **slurm-02**.
- The Scaler program can run on only one node.

Software Deployment

- Installation node

  Master node of the Slurm cluster

- Installation directory

  /data

- Downloading software

  cd /data

  Download **scaler-0.0.1-SNAPSHOT.jar**.

- Creating a configuration file

  Create the **scalerConfig.yaml** file in the directory that saves **scaler-0.0.1-SNAPSHOT.jar**, and edit the file. (Ensure that the file complies with the YAML file specifications.)

  Configure the file as follows:

```
user:
  # AK of the account for logging in to the console
  ak:
  # SK of the account for logging in to the console
  sk:
 # Tenant ID of the target region
  project: cc515cbccbc04b78b29a30f5c47fc99a
  # Proxy address, port, username, and password. They are not required if no proxy is needed.
  proxy-address:
  proxy-port:
  proxy-username:
  proxy-password:
as:
  # AS endpoint corresponding to the target region
  endpoint: as.cn-north-9.myhuaweicloud.com
  # ID of the preconfigured AS group
  group: de2aa26c-12f7-4882-824f-6c2886bb91e1
  # Maximum number of instances displayed on a page. The default value is 100. You do not need to change the value.
  list-instance-limit: 100
 # Maximum number of instances that can be deleted. This maximum number in AS is 50. You do not need to change the value.
  delete-instance-limit: 50
ecs:
  # ECS endpoint corresponding to the target region
  endpoint: ecs.cn-north-9.myhuaweicloud.com
metric:
  # Namespace of the custom monitoring metric. You do not need to change the value.
  namespace: HPC.SLURM
  # Name of a custom metric.
  name: workload
 # Name of a custom metric dimension. You do not need to change the value.
  dimension-name: hpcslurm01
  # ID of a custom metric dimension. It can be set to the AS group ID. This value does not affect
```

 Copyright © Huawei Technologies Co., Ltd.

```
functions.
  dimension-id: de2aa26c-12f7-4882-824f-6c2886bb91e1
  # Time to live (TTL) reported by the metric. You do not need to change the value.
  report-ttl: 172800
  # Cloud Eye endpoint corresponding to the target region
  endpoint: ces.cn-north-9.myhuaweicloud.com
task:
  # Period for checking the Slurm node status, in seconds
  health-audit-period: 30
  # Period for reporting custom metrics, in seconds
  metric-report-period: 30
  # Period for checking whether scale-in is required, in seconds
  scale-in-period: 120
  # Period for automatically deleting nodes, in seconds
  delete-instance-period: 10
 # Period for automatically discovering new nodes
  discover-instance-period: 20
  # Period for comparing the number of nodes in the AS group with that of nodes in the Slurm cluster, in seconds
  diff-instance-and-node-period: 60
slurm:
  # Name of the stable node. Separate multiple names with commas (,).
  stable-nodes: node-ind07urq,node-or7qht4s
  # Slurm partition where stable nodes are located
  stable-partition: dyn1
  # Slurm partition where variable nodes are located
  variable-partition: dyn1
  # Period of idle time after which idle nodes will be deleted, in seconds
  scale-in-time: 300
  # Job waiting time. If the waiting time of a job is longer than this value, the job is considered to be in a queue and will be counted in related metrics. The recommended value is 0.
  job-wait-time: 0
 # Timeout interval for registering a new node with Slurm. If a new node fails to be registered after the timeout interval expires, the node will be deleted by AS. The recommended value is 10 minutes.
  register-timeout-minutes: 10
  # Number of CPU cores of the elastic node
  cpu: 4
  # Memory size of the elastic node. This is a reserved field. It can be set to any value greater than 0.
  memory: 12600
```

You should change the values of the parameters in bold as needed. Retain the default values for other parameters.

- Installing JDK

  If JDK 8 is not installed on the master node, manually install it.

Software Management

- Starting the software

  Run the **nohup java -jar scaler-0.0.1-SNAPSHOT.jar --spring.config.name=scalerConfig > /dev/null 2>&1 &** command as the **root** user.

- Stopping the software

  a. Run the **ps –ef | grep java | grep scaler** command to query the process ID of the software.

  b. Run the **kill -9** *Process ID* command.

- Updating the configuration file

  After updating the **scalerConfig.yaml** configuration file, restart Scaler to apply the modification.

 Copyright © Huawei Technologies Co., Ltd.

Troubleshooting

Log printing warning

```
node: [Node name] status isn't DOWN
```

This message indicates that the node is not in the AS group and is not in the down state.

Check whether the node is a stable node. If yes, add the node to the **scalerConfig.yaml** file.

Relevant Files

scaler-0.0.1-SNAPSHOT.jar: main software of Scaler

scalerConfig.yaml: configuration file

scaler.log/scaler.yyyy-mm-dd.log: program run log

3.6.8 Bin Packing

Bin packing schedules multiple jobs to nodes in a compact manner to maximize resource utilization and reduce resource fragments. Slurm uses the priority-based scheduling algorithm by default. However, you can configure and adjust parameters to implement bin packing.

Bin Packing of the Slurm Scheduler

1. Resource allocation policies
   - **Default policy**: Slurm uses the **First Fit** policy (to allocate resources by node) by default. You can change the policy to **Best Fit** or **Consumable Resource (CR) optimization**.
   - **Enable compact allocation**.
     ```
     # Enable compact allocation in the slurm.conf file.
     SelectType=select/cons_tres
     SelectTypeParameters=CR_CPU_Memory
     ```
     This setting enables the scheduler to preferentially allocate CPU and memory in a compact manner during resource allocation.
2. **Node sharing mode**
   - **Node sharing**: Multiple jobs can share node resources (with CPU-based allocation being the typical configuration).
     ```
     # Configure node sharing in the slurm.conf file.
     NodeName=node[1-100] Sockets=1 CoresPerSocket=16 ThreadsPerCore=2 RealMemory=64000
     ```
   - **Job parameters**:
     ```
     sbatch -p [partition-name] --ntasks=4 --cpus-per-task=4 --mem=4G job.sh # Apply for 4 vCPUs and 4 GB of memory for each task.
     ```
3. **Forcible compact allocation**

   Use the **--distribution** parameter to forcibly pack tasks as closely together as possible within nodes.

   ```
   sbatch -p [partition] --nodes=2 --ntasks-per-node=8 --distribution=block:block job.sh
   ```

Configuration Examples

Example 1: Restricting bin packing by QoS and partition

 Copyright © Huawei Technologies Co., Ltd.

1. **Creating a QoS to limit the resources of a single job**
   ```
   sacctmgr add qos packed_qos MaxCPUsPerUser=100 MaxMemPerUser=100G
   ```
2. **Specifying resources when submitting a job**
   ```
   sbatch --qos=packed_qos --ntasks=10 --cpus-per-task=2 --mem-per-cpu=2G job.sh
   ```

Example 2: Using --exclusive and sharing together

- **Exclusive node** (This avoids resource competition but may waste resources.)
  ```
  sbatch --exclusive --nodes=1 job.sh
  ```
- **Shared node** (This improves the resource utilization.)
  ```
  sbatch --share --ntasks=4 job.sh # Allow other jobs to share the remaining resources of the node.
  ```

Monitoring and debugging

1. **Viewing node resource usage**
   ```
   sinfo -N -o "%N %C %e %m" # Displays the CPU and memory usages of the node.
   ```
2. **Analyzing job distribution**
   ```
   squeue -o "%i %P %u %T %C %m %N" # View the CPU/memory requirements and the nodes that are allocated to a job.
   ```
3. **Detecting resource fragments**
   ```
   scontrol show nodes # Check for the idle resources of the node.
   ```

Precautions

1. **Resource oversubscription risks**: Excessive compactness may cause resource contention (for example, insufficient memory). You need to monitor OOM events.
2. **Impact on job priority**: Running high-priority jobs may interrupt compact allocation. You need to strike a balance between fair access and resource utilization.
3. **Configuration complexity**: Enabling complex scheduling policies (such as **cons_tres**) may increase the scheduling delay.

You can configure resource allocation policies, job parameters, and QoS restrictions to implement efficient bin packing operations in the cockpit and significantly improve cluster resource utilization.

 Copyright © Huawei Technologies Co., Ltd.

4 FAQs

4.1 Why Am I Still Charged After I Deleted My Cluster?

- You deselected **Delete cloud server** when you delete a pay-per-use cluster. Therefore, only the cluster is deleted, but cloud servers in the cluster are still being billed.
- When you delete a yearly/monthly cluster, only the cluster is deleted, but cloud servers in the cluster are still being billed. To stop billing, manually unsubscribe from the cloud servers.
- You have purchased an EIP for the master node in the cluster, but the EIP is not deleted even if you select **Delete cloud server** when deleting the cluster. To stop billing, manually delete the EIP.
- If an EVS disk is attached to a node in the cluster, the disk is not deleted and is still being billed even if you select **Delete cloud server** when you delete the cluster. To stop billing, manually delete the EVS disk.

4.2 What Can I Do If a Compute Node Cannot Be Found or Fails to Be Added to a Cluster?

Symptom

If a compute node cannot be found or fails to be added to a cluster, the cluster cannot be deployed or created.

Handling Methods

- Delete the cluster. Manually bind an EIP to a compute node that was successfully created and then use it together with other compute nodes to create a new cluster.
- Delete the abnormal compute node to bring the cluster back to work.

 Copyright © Huawei Technologies Co., Ltd.

4.3 If I Fail to Deploy a Cloud Server in a Cluster, Can I Deploy It in Another Cluster?

Deselect **Delete cloud server** and delete this cluster. Use the cloud server to deploy another cluster.

4.4 What Can I Do If I Receive Error Message "Insufficient EIP quota"?

Symptom

Error message **Insufficient EIP quota** is displayed after I submit a request for creating a cluster.

Handling Methods

Take one of the following actions:

- Set **EIP** to **Use existing**.
- Submit a service ticket to increase your EIP quota.

4.5 What Can I Do If the Maximum Number of Clusters Has Been Reached?

Each user can create up to three clusters. You can delete the clusters that are no longer needed and create new ones again.

4.6 What Can I Do If I Receive Error Message "master node cannot reached" But the Master Node is Running?

Symptom

The master node is **Running**, but error message **master node cannot reached** is displayed.

Handling Methods

1. Check that **/usr/local/.hpccluster/role/.ssh/authorized_keys** has been added to **AuthorizedKeysFile** in **/etc/ssh/sshd_config**.
2. Check that **authorized_keys** exists in **/usr/local/.hpccluster/role/.ssh/**.

 Copyright © Huawei Technologies Co., Ltd.