DataArtsFabric

FAQ

Issue 01

Date 2025-08-25

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: https://www.huaweicloud.com/intl/en-us/

Contents

1 What Are Regions and AZs?

Regions and availability zones (AZs) are typically used to describe the locations of data centers. You can create resources in specific regions and AZs.

- A region is a physical data center location. Each region operates independently, enabling maximum fault tolerance and stability. Once a resource is successfully created, its region cannot be changed.
- An AZ is a physical location where resources use independent power supplies and networks. Within a single region, there can be multiple AZs that are physically isolated from each other yet interconnected through internal networks. This ensures the independence of each AZ while offering cost-effective, low-latency network connectivity.

Figure 1-1 Regions and AZs

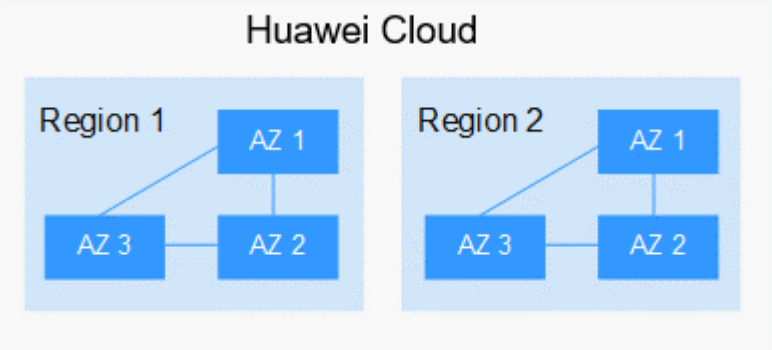

Currently, Huawei Cloud has provided cloud services in multiple regions worldwide. You can choose the region and AZ that best suits your needs. For more information, see **Huawei Cloud Global Regions**.

2 What Is a Project?

Each region in the cloud corresponds to a project by default. These projects isolate compute, storage, and networking resources across physical regions. Authorization is managed at the regional level through these default projects, allowing IAM users to access all resources within that particular region under your account.

For more refined access control, create sub-projects under a default project, purchase resources in the sub-projects, and grant resource permissions to IAM users.

NOTE

Resources in a project cannot be transferred.

3 What Is Ray?

Concepts

Ray is an open-source framework for building and running distributed applications. It was developed by the University of California, Berkeley, RISELab. It aims to simplify distributed computing of machine learning (ML) workloads and supports complex task graphs. This enables developers to write parallel or distributed programs that can efficiently use all available resources in the cluster.

Ray provides an easy-to-use API to implement parallel and distributed execution of tasks. It allows users to convert Python functions into remote procedure calls (RPCs) so that they can be concurrently executed on multiple nodes. In addition, Ray introduces the concept of dynamic task graph, which enables it to handle workloads that require flexible scheduling, such as reinforcement learning, hyperparameter tuning, and other iterative algorithms.

By providing support for distributed computing, Ray promotes faster model training and more efficient resource usage, and is a powerful tool for researchers and engineers who want to expand their applications on multiple machines. Additionally, the Ray ecosystem includes some advanced libraries, such as Ray Tune (for hyperparameter adjustment), RLlib (for reinforcement learning), and Ray Serve (for model services), to meet requirements in different scenarios.

Use Cases

As a distributed machine learning computing framework, Ray is typically used in scenarios such as model training and model fine-tuning. The parallel computing capability provided by Ray can significantly improve the computing efficiency.

