Multi-Cloud Container Platform

Service Overview

Issue 01

Date 2022-03-29

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.

Contents

1 What Is Multi-Cloud Container Platform?

Multi-Cloud Container Platform (MCP) is developed by HUAWEI CLOUD based on years of experience in the cloud container field and a community-advanced cluster federation technology (**Karmada**). It provides multi-cloud and hybrid cloud solutions for unified cluster management across clouds and unified deployment and traffic distribution of applications across clusters. It not only resolves multi-cloud disaster recovery, but also plays an important role in traffic sharing, decoupling of data storage and service processing, decoupling of development and production environments, and flexible allocation of computing resources.

Figure 1-1 Using MCP

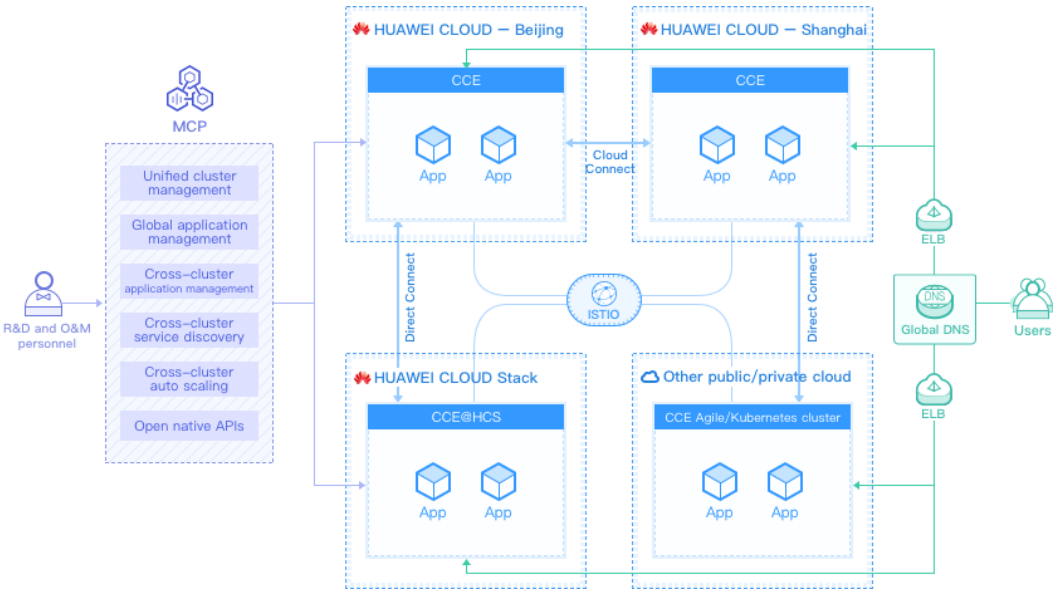

Function Description

- **Unified cluster management**

  MCP leverages cluster federation to implement unified management of clusters of different cloud service providers. As a unified entry for multiple

clusters, MCP supports dynamic cluster access and global cluster monitoring dashboard, which facilitates cluster deployment, release, and O&M.

- **Global application management**

  MCP manages Kubernetes clusters in different regions or clouds and supports full lifecycle management of applications across clusters, including deployment, deletion, and upgrade, by using standard cluster federation APIs in Kubernetes.

- **Cross-cluster auto scaling**

  MCP supports cross-cluster auto scaling policies to balance the pod distribution in each cluster and implement global load balancing. You do not need to worry about horizontal scaling of cluster nodes. MCP automatically scales in or out resources required by applications based on the application load.

- **Cross-cluster service discovery**

  MCP supports **federated service** creation and cross-cluster service discovery. It implements service region affinity based on the proximity access principle. In this way, when a service is deployed in multiple regions, the one in the region where the user is located is preferentially accessed, reducing network latency.

- **Standards compliance**

  MCP is compatible with the latest federation architecture of the Kubernetes community and provides native Kubernetes APIs and Karmada APIs. You can use the Kubernetes command lines and APIs to deploy containerized applications without modifying any service code.

- **Application federation**

  MCP supports application federation, which allows you to deploy an application from only one cluster to multiple clusters across clouds in just a few clicks. In this way, cross-cloud DR and traffic sharing can be implemented.

- **Cross-cluster application cloning and migration**

  You can clone or migrate your applications to other clusters or across clouds/regions in just a few clicks without the need of re-writing or modifying your service code.

Karmada

Karmada is a Kubernetes API-based system for multi-cluster management. In multi-cloud and hybrid cloud scenarios, Karmada provides ready-to-use functions and automated managed multi-cluster applications to implement centralized multi-cloud management, high availability, fault recovery, and traffic scheduling.

Key features

- Cross-cluster application management based on Kubernetes APIs allows you to quickly deploy applications from a single cluster to multiple clusters.
- Kubernetes clusters are operated and managed in a centralized manner.
- Cross-cluster applications support automatic extension, failover, and load balancing across multiple clusters.
- Advanced scheduling policies are available. Applications can be scheduled based on the region, availability zone, cloud provider, and cluster affinity/anti-affinity.

- Creation and distribution of CustomResourceDefinition (CRD) resources are supported.

Architecture

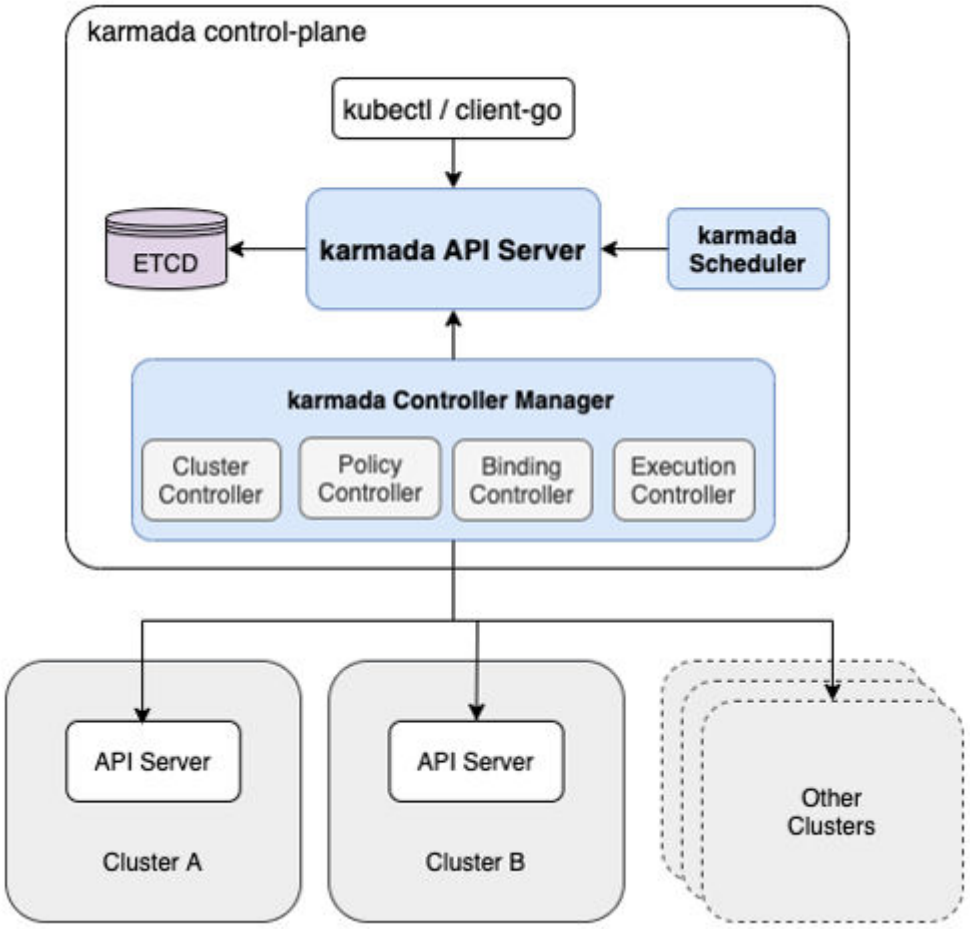

- ETCD: Stores Karmada API objects.
- Karmada Scheduler: Provides advanced multi-cluster scheduling policies.
- Karmada Controller Manager: Contains multiple controllers, which listen to the Karmada objects and communicate with the API servers on the member clusters.

Basic concepts

- Resource Template: Karmada uses native Kubernetes API definitions as resource templates to quickly interconnect with the Kubernetes ecosystem tool chain.
- Propagation Policy: Karmada provides independent policy APIs to configure resource propagation policies.
- Override Policy: Karmada provides independent differentiated APIs to configure differentiated configurations related to the clusters. For example, you can use Override Policy to configure different clusters to use different images.

The following figure shows the Karmada resource propagation process:

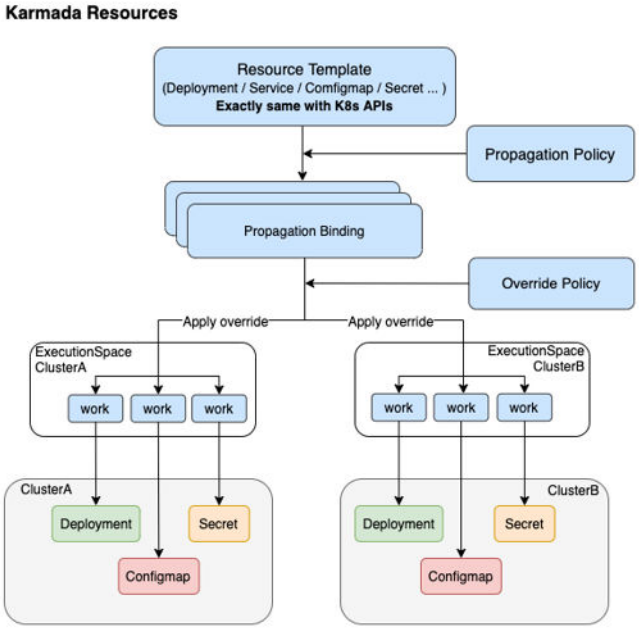

2 Advantages

In the cloud computing era, a large number of enterprises host their services on the clouds. However, cloud technologies have their own limits, and CIOs are trying their best to avoid cloud service failures. One of the most effective solutions is to deploy services on multiple clouds (including public and private clouds) for disaster recovery. This solution can address the following requirements of enterprises in single-cloud scenarios:

- Avoid vendor lock-in and enhance enterprises' business flexibility.
- Provide multi-cloud DR to quickly cope with single point of failure (SPOF) in the cloud.
- Intelligently balance service traffic loads to cope with the impact of traffic spikes.
- Deploy services across regions, achieving region-based service access and improving performance.

Among various multi-cloud/hybrid-cloud solutions, the combination of multi-cloud/hybrid-cloud and containers is a better choice. HUAWEI CLOUD MCP is developed to address the preceding requirements. Compared with traditional multi & hybrid cloud solutions, MCP has the following advantages:

- Unified container technology standard: Applications can be flexibly migrated among multiple clusters across clouds, without worrying about the dependency on the environment.
- Scaling within seconds: Enterprises do not need to maintain additional resources for multi & hybrid cloud solutions and therefore extra investment is not required.
- Lightweight technology solution: Enterprises can build and maintain their cloud services easily, without paying attention to a large number of infrastructure problems.

Networking

MCP provides three types of networking solutions: hybrid-cloud, multi-cloud, and multi-region. You can add HUAWEI CLOUD CCE clusters, Kubernetes clusters of other public clouds (such as Alibaba Cloud and Tencent Cloud), and on-premises Kubernetes clusters to MCP for unified management.

Hybrid cloud solution: MCP manages HUAWEI CLOUD CCE clusters and Kubernetes clusters in on-premises data centers in a unified manner.

This solution has the following characteristics:

- Unified management on the cloud: Resources, applications, and networks are managed on the public cloud in a unified manner.
- Application scaling on the cloud: The private cloud clusters are used as fixed resources, and the public cloud clusters are used as backups. When a traffic spike occurs, a part of the traffic can be quickly redirected to cloud clusters.
- Decoupling of data storage and service processing: Core data is stored in on-premises clusters, and applications run on cloud clusters.
- Smooth service migration: Services are gradually migrated to the public cloud to maximize the utilization of physical computing resources.

Figure 2-1 How the hybrid cloud solution works

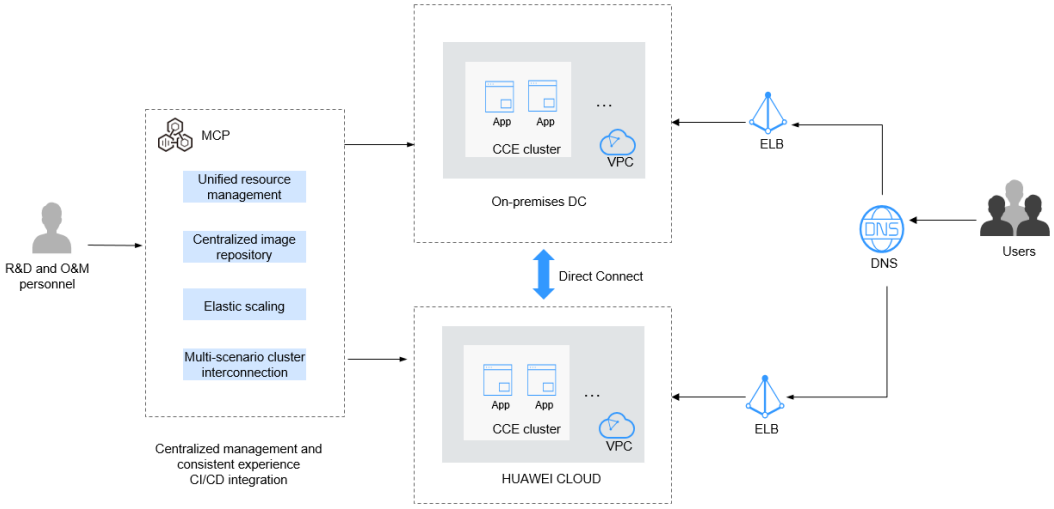

Multi-cloud solution: MCP manages HUAWEI CLOUD CCE clusters and Kubernetes clusters of other cloud service providers (such as Alibaba Cloud and Tencent Cloud) in a unified manner.

This solution has the following characteristics:

- Unified management: Resources, applications, and networks are managed in a unified manner.
- Scalability: Higher horizontal scalability provides more capacity than a single cluster.
- Intelligent routing: Clusters are distributed in different regions, and requests are forwarded from a geographically closer cloud, which in turn reduces network latency caused by cross-region access.
- Multi-cloud DR: The same applications are deployed in clusters of different cloud service providers to implement disaster recovery.

Figure 2-2 How the multi-cloud solution works

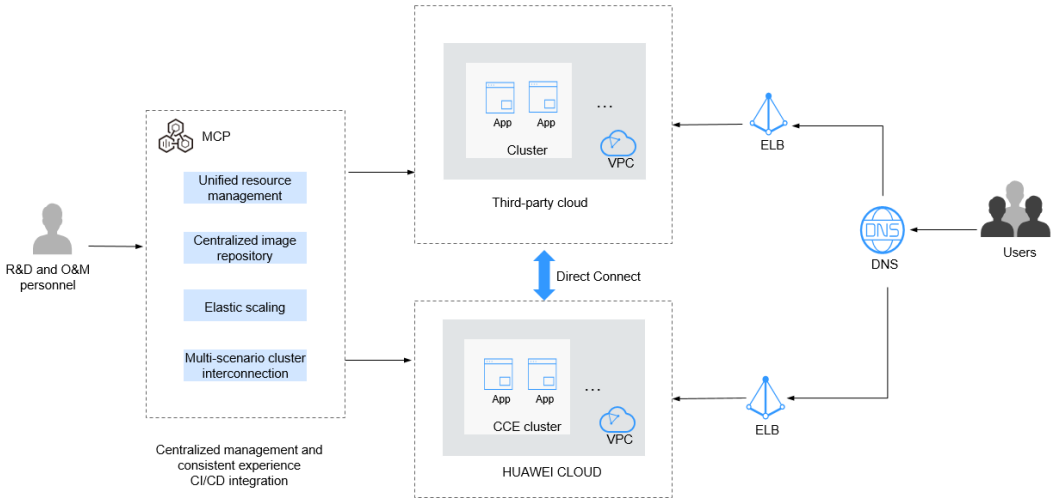

Multi-region solution: MCP manages CCE clusters in the same or different VPCs or regions of HUAWEI CLOUD in a unified manner.

This solution has the following characteristics:

- Unified management: Resources, applications, and networks are managed in a unified manner.
- Scalability: Higher horizontal scalability provides more capacity than a single cluster.
- Intelligent routing: Clusters are distributed in different regions, and requests are forwarded from a geographically closer cloud, which in turn reduces network latency caused by cross-region access.
- Multi-region DR: The same applications are deployed in clusters across different regions of HUAWEI CLOUD to implement disaster recovery.

Figure 2-3 How the multi-region solution works

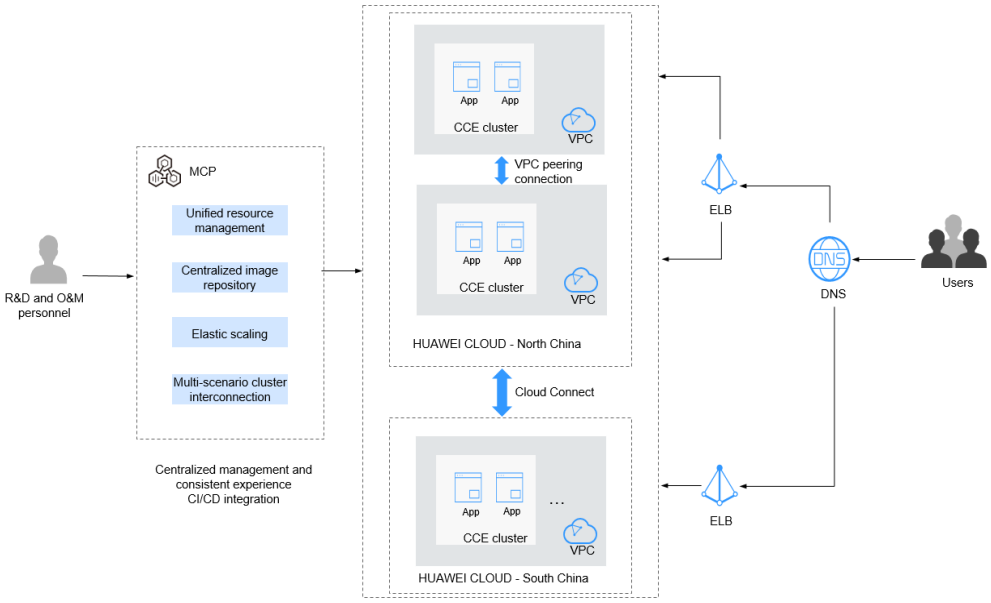

3 Application Scenarios

Traffic Sharing

In industries such as e-commerce, video on-live, and education, traffic surges in a specific period due to scenarios such as flash sales, live interaction, score query, or registration. To cope with the impact of traffic spikes and improve service stability, MCP provides cross-cloud auto scaling based on traffic policies. With cross-cloud and cross-cluster scaling for application instances and unified management of application traffic, MCP implements elastic load balancing of traffic across different clouds. You can configure a traffic policy when deploying an application across clouds. If the private cloud clusters or clusters of a public cloud are overloaded due to traffic spikes, MCP flexibly schedules services to other public cloud clusters based on the traffic policy. This policy can prevent the system from breaking down due to traffic bursts.

- **Advantages**
  - Scalability: Higher horizontal scalability provides more capacity than a single cluster.
  - Intelligent routing: Clusters are distributed in different regions, and requests are forwarded from a geographically closer cloud, which in turn reduces latency.
  - Auto scaling: Application instances are automatically scheduled to different clouds based on the configured scaling policies.
- Recommended Services for Use

  Cloud Container Engine (CCE), Software Repository for Container (SWR), and Elastic Load Balance (ELB)

Figure 3-1 How traffic sharing works

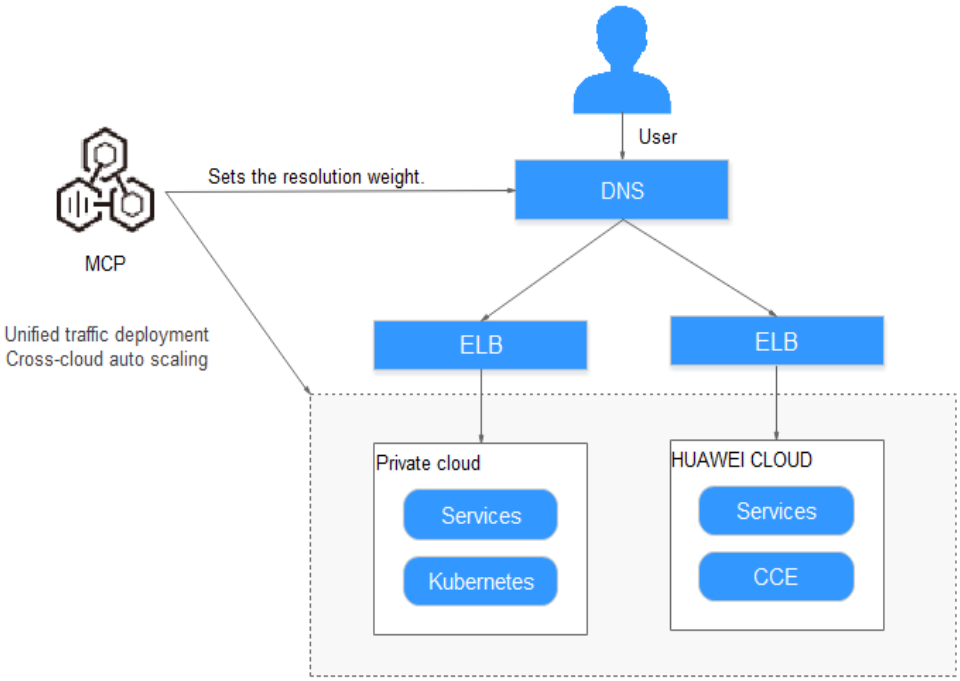

Cross-Cloud DR

To cope with failures of a single cloud, MCP allows instances of an application to run on multiple clouds. When one of the clouds is down, MCP will migrate instances and switch over traffic to the other clouds within seconds. In this way, service reliability is greatly improved. In addition, compared with traditional multi & hybrid cloud solutions, containers can be scaled within seconds, eliminating the need to maintain redundant resources for DR and reducing the construction and O&M costs of infrastructure resources.

- **Advantages**
  - Unified multi-cluster management to automatically monitor the health status of each cluster.
  - Unified application management, auto scaling, and automatic migration of application instances to other clouds during DR.
  - Unified traffic management and automatic traffic switchover.
- **Recommended Services for Use**

  CCE, SWR, ELB, and Domain Name Service (DNS)

Figure 3-2 How cross-cloud DR works

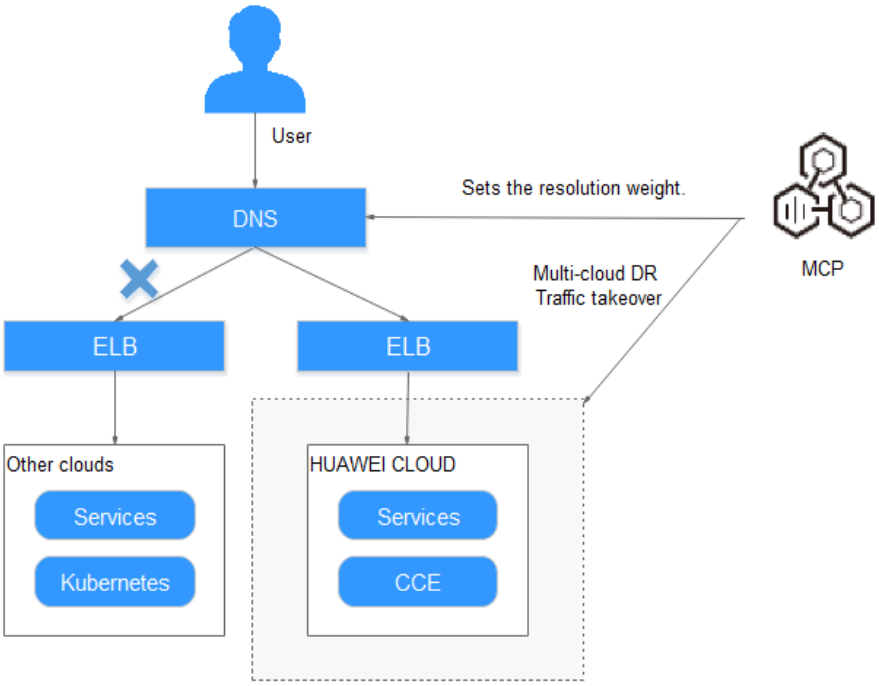

Decoupling of Data Storage and Service Processing

For security purpose, users in the finance and security industries require that core services run on their private cloud clusters. With HUAWEI CLOUD MCP, you can store sensitive data in the private cloud clusters and run common services on the public cloud to ensure data security.

- **Advantages**
  - Sensitive data is completely independent and controllable, eliminating information security risks.
  - Unified management of service resources reduces O&M workload.
  - Networks are connected through Direct Connect (DC), which features high performance and reliability.
- **Recommended Services for Use**

  CCE, SWR, ELB, and DC

Figure 3-3 Decoupling of data storage and service processing

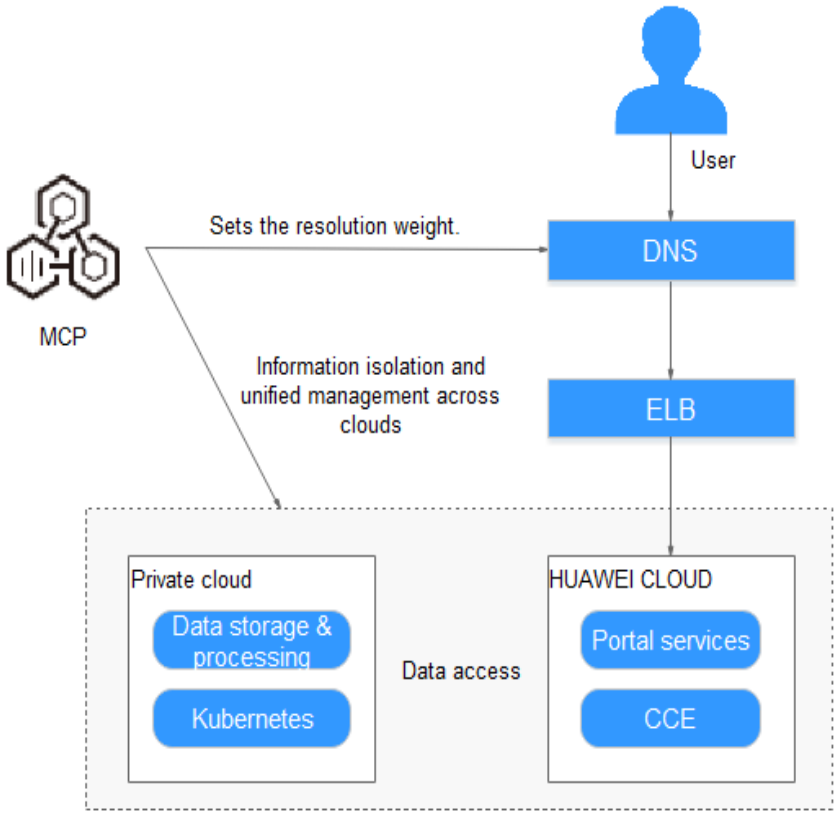

Decoupling of Development and Production Environments

For network security in continuous integration (CI) or continuous delivery (CD) scenarios, some users want to deploy development and test environments in the private cloud clusters and the production environment in the public cloud clusters. With HUAWEI CLOUD MCP, you can manage clusters where the development, test, and production environments run in a unified manner. MCP works with ContainerOps to implement service rollout by pipeline, improving the code delivery and deployment efficiency.

- **Advantages**
  - Multiple environments are decoupled.
  - The consistent running environment eliminates environment dependency for service rollout.
  - ContainerOps implements process automation from code to service rollout.
- **Recommended Services for Use**

  CCE and SWR

Figure 3-4 Decoupling of development and production environments

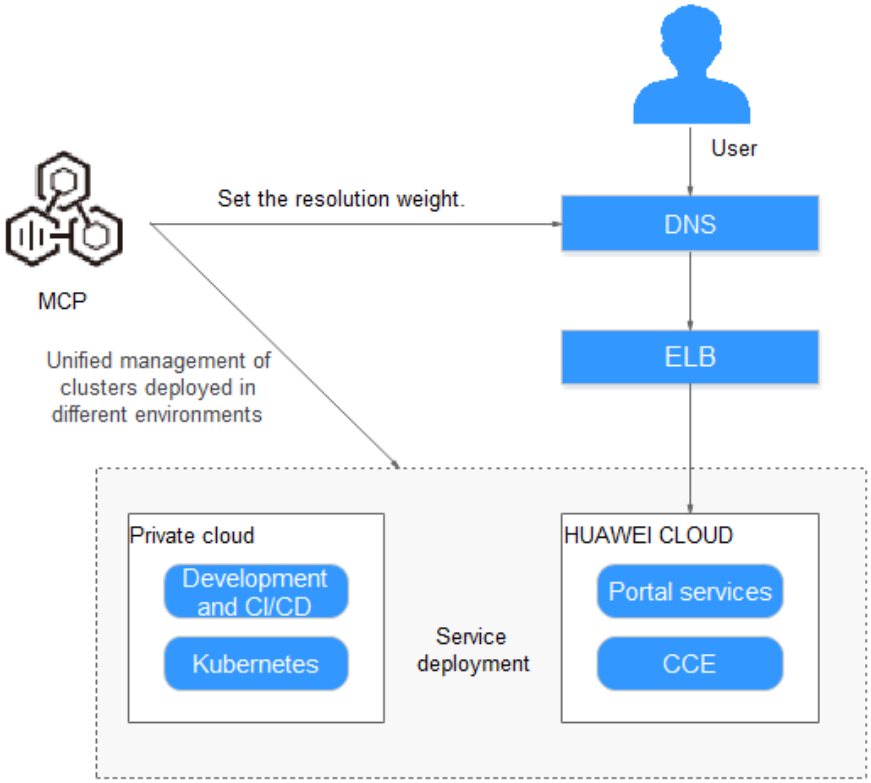

Flexible Allocation of Computing Resources

Computing tasks in industries such as AI, genomic sequencing, and video processing rely on GPUs, bare metal servers, and other special hardware. With MCP, you can deploy compute-intensive applications in the cloud and common applications in the private cloud or other clouds, avoiding high costs caused by large-scale use of special hardware.

- **Advantages**
  - Hardware leasing greatly reduces the O&M costs of physical resources.
  - Rapid scaling reduces idle resources and supports on-demand procurement.
  - Huawei-developed elastic bare metal servers are supported to deliver optimal performance for customers.
- **Recommended Services for Use**

  CCE, Bare Metal Server (BMS), and Object Storage Service (OBS)

Figure 3-5 Flexible allocation of computing resources

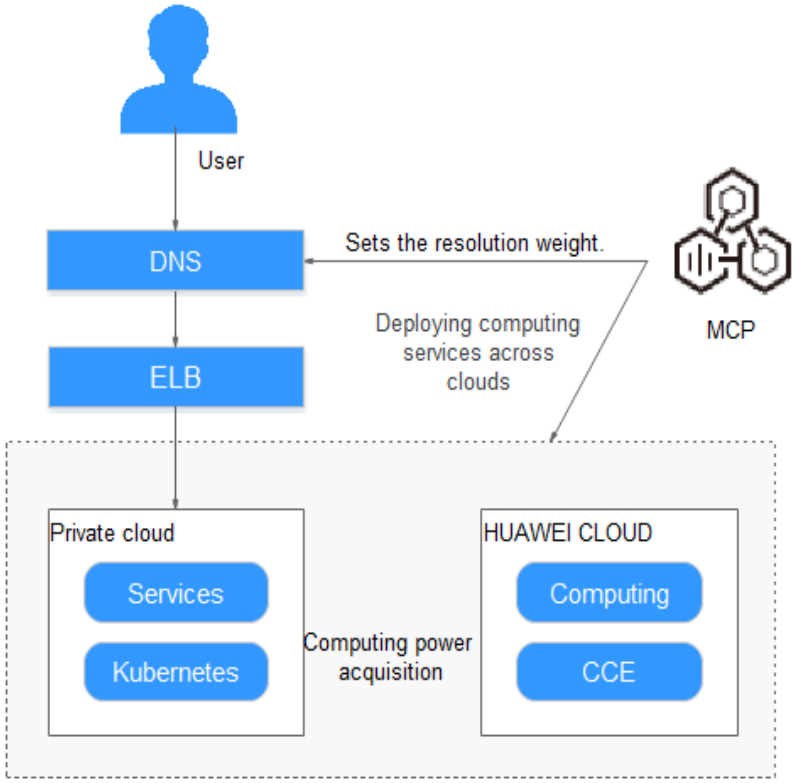

4 Constraints

Kubernetes Version Constraints

Kubernetes clusters connected to MCP must be v1.11 or later.

Other Constraints

When federated resources (such as federated Deployments, services, and PVCs) are created on MCP, the corresponding Kubernetes resources (such as Deployments, services, and PVCs) will be automatically created in the clusters connected to MCP. You must provide the **admin** account credential for accessing the clusters.

5 Permission Management

This chapter describes how to assign different permissions to employees in your enterprise to access your MCP resources. If your Huawei Cloud account does not need to create individual access credentials and permissions for your employees, skip this chapter.

You can use Identity and Access Management (IAM) to create individual **IAM users** under your account for your employees and assign them permissions to control their access to Huawei Cloud resources. Employees can use their IAM usernames and passwords to log in to Huawei Cloud and use resources based on assigned permissions.

IAM is the basic permission management service of Huawei Cloud and can be used free of charge. For more information about IAM, see **IAM Service Overview**.

MCP Permissions

By default, new IAM users do not have any permissions. After the administrator assigns permissions to them on the IAM console, they can perform specified operations on cloud services.

NOTICE

Before creating an IAM user to use MCP, **authorize access for MCP** in the current region using your Huawei Cloud account.

Table 5-1 lists all the system roles on which MCP depends. You need to assign these roles of other Huawei Cloud services to IAM users before they can use MCP.

Table 5-1 System roles on which MCP depends

Role Name	Description	Dependency
CCE Administrator	Read and write permissions for CCE clusters and all resources (including workloads, nodes, jobs, and Services) in the clusters.	Users granted permissions of this role must also be granted permissions of the following roles/policies: Global service project: OBS Buckets Viewer and OBS Administrator Region-specific projects: Tenant Guest, Server Administrator, ELB Administrator, SFS Administrator, SWR Admin, and APM FullAccess
VPC Administrator	Permissions for VPC, excluding permissions for creating, modifying, deleting, and viewing security groups and security group rules.	It depends on the **Tenant Guest** role. Users granted permissions of the **VPC Administrator** role must also be granted permissions of the **Tenant Guest** role.

6 Basic Concepts

For details about the basic concepts of Kubernetes clusters, see **Basic Concepts of CCE**.

Cluster

A cluster is a group of one or more cloud servers (also known as nodes) in the same subnet. A cluster can be seen as one or more elastic cloud servers (also called nodes) in a same subnet. It provides computing resource pool for the container running through computer groups formed by relevant technologies.

Pod

A pod is the smallest and simplest unit in the Kubernetes object model that you create or deploy. A pod encapsulates an application container (or, in some cases, multiple containers), storage resources, a unique network IP address, and options that govern how the container(s) should run.

Workload

A workload is an abstract model of a group of pods in Kubernetes. Workloads defined in Kubernetes include Deployments, StatefulSets, jobs, and DaemonSets.

- **Deployment**: Pods of a deployment are completely independent of each other and deliver the same functions. Deployments support auto scaling and rolling updates. Typical examples include Nginx and WordPress.
- **StatefulSet**: Pods in a StatefulSet are not completely independent of each other. StatefulSets have stable persistent storage and stable unique network identifiers. They support ordered, graceful deployment, scaling, and deletion. Typical examples include MySQL-HA and etcd.
- **Job**: It is a one-time task that runs to completion. It can be executed immediately after being created.
- **Cron job**: It runs a job periodically on a given schedule. For example, you can create a cron job that will perform time synchronization for all active nodes at a fixed time point.

Figure 6-1 Relationship between workloads and pods

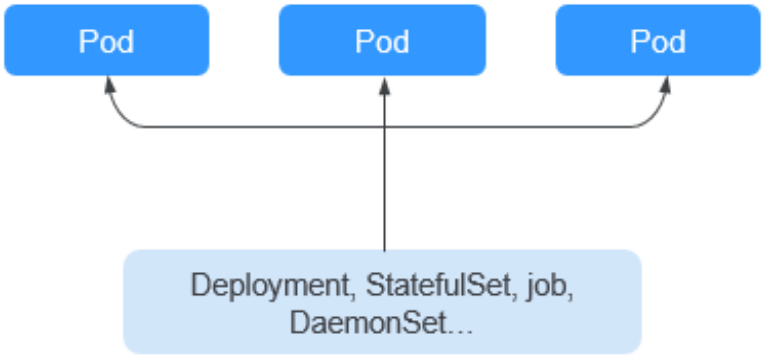

Federated Workload

A federated workload describes a workload deployed across clusters. You can use federated workloads to deploy workloads to Kubernetes clusters you selected. Currently, only federated Deployments are supported.

Federated Service

Different workload access types are provided to address diverse scenarios. In multi-cloud scenarios, a federated service is added for mutual access between workloads in different clusters in the federation. You can manage these services on MCP in a centralized manner.

Federated Ingress

Enhanced ELB is used for an ingress. Compared with layer-4 ELB, layer-7 ELB newly supports Uniform Resource Identifiers (URI) configurations and distributes access traffic to the corresponding service based on the corresponding URIs. A federated ingress refers to the ingress component used in multi-cloud scenarios. The federated ingresses can be managed on MCP in a unified manner.

Federated PVC

A PersistentVolumeClaim (PVC) refers to cluster file storage. However, a federated PVC is oriented to the federation layer. File storage is created for multiple clusters at the same time, and differentiated settings can be configured for each cluster. A federated PVC manages file storage in clusters in a unified manner.

Component

A component is the minimum unit in MCP. It provides an independent function module or microservice and supports unified O&M in multi-cloud scenarios.

ContainerOps

ContainerOps enables continuous integration (CI) and continuous delivery (CD) of containerized applications throughout the entire process from source code to deployment, providing capabilities such as image repository, image build, version management, and release pipelines.

7 Related Services

CCE

Cloud Container Engine (CCE) is a high-performance, scalable container service through which enterprises can build high-reliability containerized applications. With integrated HUAWEI CLOUD network and storage capabilities and support for Kubernetes and Docker, CCE makes it simple to build and manage diverse containerized applications. It also provides various efficient O&M functions such as container fault self-healing, monitoring log collection, and auto scaling.

CCE clusters and other Kubernetes clusters can be added to MCP for unified management.

ELB

Elastic Load Balance (ELB) automatically distributes access traffic to multiple Elastic Cloud Servers (ECSs) to balance the loads. It enhances an application's fault tolerance level and capabilities of providing services.

MCP supports cross-cluster auto scaling policies to balance the pod distribution in each cluster and implement global load balancing. You do not need to worry about horizontal scaling of cluster nodes. MCP automatically scales in or out resources required by applications based on the application load.

You can access a container workload from an external network through an elastic load balancer.

