

Cloud Application Engine

Service Overview

Issue	01
Date	2025-08-27



Copyright © Huawei Cloud Computing 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 Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions



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

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

Notice

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

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

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 Is CAE?.....

2 Advantages.....

3 Typical Scenarios.....

3.1 Microservice Application Management.....

3.2 Web Application Lifecycle Management.....

4 Glossary.....

5 Specifications.....

6 Permissions Management.....

7 Relationship with Other Cloud Services.....

1

3

5

5

6

7

9

12

20

1 What Is CAE?

Cloud Application Engine (CAE) is serverless hosting for ultra-fast web and microservice application deployment. This hosting is low cost and zero infrastructure O&M. CAE releases applications from source code, software packages, and container images quickly and easily, with auto scaling to the second, all billed pay-per-use. The whole application lifecycle is manageable with observable metrics.

Component Management and Configuration

Once developed, application components can be hosted on the CAE for:

Complete lifecycle management

- Creation and deployment from source code, software packages (JAR, WAR, or ZIP), and container images
- Control of the entire process from application component creation to logout—covering component creation, deployment, start, upgrade, rollback, and stop

Automatic component expansion without O&M

- Using VM resources is fine-grained billed to the number of requests. Application instances scale to user traffic for full cloud hosting of user applications.
- Underlying IaaS simplifies cloud migration since users no longer need to carry out O&M nor learn about complex underlying resources.

Component Monitoring and Logs

To quickly detect performance bottlenecks and locate causes, the problem (such as slow page loading or frame freezing) needs to be reproducible. Troubleshooting examples:

- When a user reports that page loading is slow, is this problem caused by a network fault, resource loading, or Document Object Model (DOM) parsing?
- Is the slow loading due to province or country where the user is located, or the user's browser or device?
- Another example, when a user reports frame freezing, is this problem caused by a faulty network between the user terminal and the server? Or is the server or database overloaded?

It is also difficult to quickly locate root causes in code. CAE addresses these difficulties by providing:

- Multi-dimensional instance monitoring for application components to give you the real-time running status of online components.
- UI-based log and event search and alarms are provided to rapidly notify you of issues.

2 Advantages

CAE is a product of Huawei's success in cloud transformation and technological innovation. This one-stop application cloud platform has the following advantages over traditional platforms:

Table 2-1 Advantages

Application Lifecycle	Traditional Hosting	CAE Hosting
Environment preparation	<ul style="list-style-type: none">• Inefficient obtaining of resources (> 1 day)• Low resource utilization (< 30%)	<ul style="list-style-type: none">• Self-service and efficient resource obtaining (in minutes)• Pay-per-use payment (auto scaling)
Installation and deployment	<ul style="list-style-type: none">• Siloed system• Manual deployment	Developers only need to use CAE and source code software repository for one-click automatic deployment and updates.
Component configuration	<ul style="list-style-type: none">• Complex items• All involved environments configured separately• Error-prone	The configuration file is decoupled from the environment so that the independently maintained file is used together by multiple environments.
Component upgrade	<ul style="list-style-type: none">• Patch installation• Manual upgrade• Services interrupted	Rolling upgrades even distribute services to new and old instances to avoid interrupting services.

Application Lifecycle	Traditional Hosting	CAE Hosting
Component O&M	<ul style="list-style-type: none"> • Application breakdown or crash • Slow service response • Insufficient system resources • Difficult fault locating 	<ul style="list-style-type: none"> • Real-time graphical display of monitoring metrics Monitors CPU usage, running logs, and key events in real time • Application performance analysis Supports application discovery, dependency, and KPI aggregation. • Tracing and monitoring Monitors platform, resources, and applications while analyzing microservice traces. • Microservice governance Provides SLA metrics for microservice APIs (throughput, latency, and success rate) and ensures continuous services with monitoring and governance in real time (response in seconds).

3 Typical Scenarios

3.1 Microservice Application Management

Service Scenarios

Usage Scenario

In a traditional single architecture, different service models need a unified technical solution and technical platform. Each service module cannot be reused, so any faulty module will cause the entire system to become unavailable. With the increasing complexity of enterprise services, the traditional monolithic architecture becomes more and more cumbersome and inflexible. Microservice applications solve these problems.

Benefits

Microservice-based applications divide a system into several small service components. Such components communicate with each other through lightweight protocols and decouple their lifecycle management.

Ever-growing services may encounter unexpected situations such as instantaneous and large-scale concurrent access, service errors, and intrusion. The microservice architecture fine-tunes service management and control to support service requirements.

CAE manages the full lifecycle of serverless applications. Compatible runtime environments include Java, PHP, Node.js, Docker, and Tomcat, and microservice applications such as Apache ServiceComb, Spring Cloud, and Dubbo are managed without intrusion. In addition, CAE improves cloud migration with configuration management, monitoring and O&M, and service governance.

Advantage

CAE microservice application solution:

- Supports multiple microservice frameworks, such as ServiceComb, Spring Cloud, Dubbo, and Service Mesh, and can be directly hosted on the cloud without changing service code.

- Supports multiple languages, such as Java, Node.js, PHP, and Python.
- Provides functions such as service center, configuration center, dashboard, and dark launch.
- Provides complete microservice governance policies, including fault tolerance, rate limiting, service degradation, circuit breaker, fault injection, and blacklist/whitelists. Operations are available on the UI for different service scenarios, greatly improving the availability of service governance.
- Implements mutual discovery between Spring Cloud, ServiceComb, and Java chassis.

3.2 Web Application Lifecycle Management

Service Scenarios

Usage Scenario

Web applications are in wide use among enterprise service systems, online store systems, forums, blogs, Wiki systems, and online games. Enterprise IT departments are responsible for managing their lifecycle with different technical architectures.

Benefits

A unified platform greatly reduces workload, improves efficiency, and quickly responds to complex and changing service requirements.

Advantage

CAE greatly improves development and O&M so enterprises can focus instead on service innovation. It has the following advantages:

- Deployment with a few clicks using WAR, JAR, or ZIP packages.
- One-stop O&M provides upgrade, rollback, log, monitoring, and auto scaling.
- Seamless integration with cloud services, such as CSE and RDS, and related applications.

4 Glossary

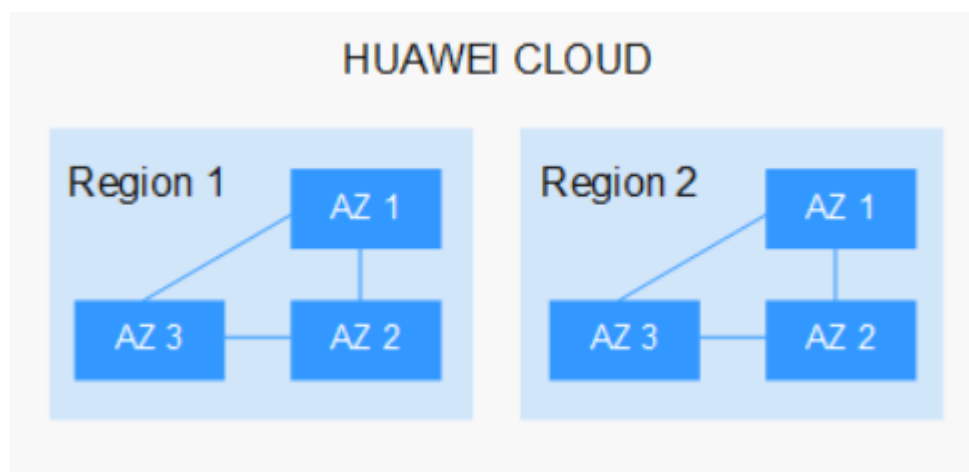
Regions and AZs

A region and availability zone (AZ) identify the location of a data center. Created resources have a specific region and AZ.

- Regions have their own geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS), are shared in the same region. Regions are either universal or dedicated: universal region provides universal cloud services for common domains while a dedicated region provides services of the same type only or for specific domains.
- An AZ contains one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. An AZ's computing, network, storage, and other resources are logically divided into multiple clusters. AZs within a region are interconnected by high-speed optical fibers for building cross-AZ high-availability systems.

Figure 4-1 shows the relationship between regions and AZs. For details about how to select regions and AZs, see [Region and AZ](#).

Figure 4-1 Regions and AZs



Environment

Environments distinguish service deployment scenarios and means isolation. In CAE, development, test, pre-production, and production environments are tailored to requirements. Networks in an environment intercommunicate, while components and services are managed and deployed within the environment for easier O&M, production, and rollout.

Application

An application is a service system with functions and consists of one or more application components.

Application Component

An application component implements a service feature of an application. It is in the form of code or software package and can be deployed independently.

Component Instance

An instance is the minimum running and deployment unit of a component, and tends to be used for a single application process.

Serverless

Serverless computing is a model of cloud computing. The cloud provider allocates machine resources on demand and manages servers on behalf of the customers, who are not involved in capacity planning, configuration, management, maintenance, fault tolerance, or scaling of containers, VMs, or physical servers.

Data Redundancy Policy

The two data redundancy policies are multi-AZ storage and single-AZ storage. CAE only supports authorization and creation of object storage with policy multi-AZ storage.

- Multi-AZ storage: Data is stored in multiple AZs in the same region for better availability. A file system created in multi-AZ mode stores its data in three different AZs in the same region. This enables data access in case of unavailability of any single AZ, so the multi-AZ mode is useful for storing data that must always be available.
- Single-AZ storage: Data is stored in a single AZ, with lower costs.

5 Specifications

This topic describes CAE specifications and restrictions.

Supported Regions

[Table 5-1](#) lists the supported regions.

Table 5-1 Supported regions

Region	Region ID
AP-Singapore	ap-southeast-3
ME-Riyadh	me-east-1
TR-Istanbul	tr-west-1
AF-Johannesburg	af-south-1
LA-Mexico City2	la-north-2

Quota Limits

A quota is the maximum number of resources that you can create. [Table 5-2](#) provides the resource quota limitations.

Table 5-2 Quota

Resource	Maximum Quota	Modifiable
Environment in each region	1	Yes
Number of components that can be created by each tenant in each region	50	Yes

Resource	Maximum Quota	Modifiable
Number of instances of each component	99	No
CPU cores	2 cores per instance	No
Memory	4 GB per instance	No

To increase environments and components, submit a [service ticket](#).

Component Specifications

[Table 5-3](#) lists memory sizes for each CPU core.

Table 5-3 Component specifications

CPU (cores)	Memory (GiB)
0.5	1
0.5	2
1	1
1	2
1	4
2	2
2	4

More specifications are coming soon.

Bandwidth Limits

[Table 5-4](#) lists bandwidth limits in public network access through EIP

Table 5-4 Bandwidth limits

Access	Bandwidth Limit (Mbit/s)
Public network to CAE	20
Application to public network	20

To increase bandwidth limit, submit a [service ticket](#).

Storage

CAE cannot read or write local disks. You are advised to mount an OBS parallel file system. For details, see [Parallel File System Overview](#).

Specifications on Temporary Storage

If you have configured a log path and selected local disk storage, CAE uses the storage to store log files. The temporary files generated by the file system or program running on the image occupy the storage space.

Containerized applications can read temporary data and write it to the temporary storage. The storage is temporary and will be cleared after the container is stopped or restarted. Files written to the storage are visible only to the processes running in the current container.

The total storage of each instance depends on the total number of CPUs allocated to the instance. For details about specifications on the temporary storage, see [Table 5-5](#). Specifications with more than three cores can be used only by whitelisted users. If you need to use such specifications, [submit a service ticket](#)

When using local disks for storage, you are advised to properly plan log storage and periodically rotate and clear logs to prevent component restart caused by excessive logs. If the log usage exceeds the limit, you are advised to use cloud storage for log configuration.

Table 5-5 Specifications on temporary storage

CPU (cores)	Total Temporary Storage (GiB)
0.5 CPUs	2 GiB
0.5 < CPUs < 4	4 GiB
4 ≤ CPUs	8 GiB

6 Permissions Management

If you need to grant your enterprise personnel permission to access your CAE resources, use Identity and Access Management (IAM). IAM provides identity authentication, permissions management, and access control, helping you secure access to your cloud resources.

With IAM, you can create IAM users and grant them permission to access only specific resources. For example, if you want some software developers in your enterprise to be able to use CAE resources but do not want them to be able to delete CAE resources or perform any other high-risk operations, you can create IAM users and grant permission to use CAE resources but not permission to delete them.

If your Huawei Cloud account does not require individual IAM users for permissions management, you can skip this section.

IAM is a free service. You only pay for the resources in your account.

For more information about IAM, see [IAM Service Overview](#).

CAE Permissions

New IAM users do not have any permissions assigned by default. You need to first add them to one or more groups and then attach policies or roles to these groups. The users then inherit permissions from the groups and can perform specified operations on cloud services based on the permissions they have been assigned.

CAE is a project-level service deployed for specific regions. When you set **Scope** to **Region-specific projects** and select the specified projects (for example, **cn-east-3**) in the specified regions (for example, **CN East-Shanghai1**), the users only have permissions for CAE in the selected projects. If you set **Scope** to **All resources**, the users have permissions for CAE in all region-specific projects. When accessing CAE, the users need to switch to the authorized region.

You can grant permissions by using roles and policies.

- **Roles:** IAM's coarse-grained authorization that defines permissions by job responsibility. This mechanism provides a limited number of service-level roles for authorization. Different services often depend on other services, so these dependencies must be considered when assigning roles. Roles are not an ideal choice for fine-grained authorization and secure access control.

- **Policies:** A fine-grained authorization tool that defines permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization, meeting requirements for secure access control.

Table 6-1 lists all the system-defined permissions for CAE.

Table 6-1 CAE system permissions

Role/Policy Name	Description	Type	Dependent System Permissions	Service Dependency and Scenario
CAE FullAccess	Full permissions for CAE.	System-defined policy	Must be used together with these permissions: <ul style="list-style-type: none">• OBS Administrator• AOM FullAccess• SWR Admin• BSS FinanceAccess• VPC ReadOnlyAccess• ELB FullAccess• ServiceStage FullAccess	Full permissions for CAE.
CAE ReadOnlyAccess	Read-only permissions for CAE.	System-defined policy	Must be used together with these permissions: <ul style="list-style-type: none">• ServiceStage ReadOnlyAccess• LTS ReadOnlyAccess• obs:bucket:GetBucketLocation• obs:bucket:GetBucketStorage• obs:bucket:GetBucketStoragePolicy	Read-only permissions for CAE.

Policies are also customizable beyond permissions listed in **Table 6-1** do not meet your requirements, as shown in **Table 6-2**.

Table 6-2 Common operations supported by each system policy

Description	CAE ReadOnlyAccess	CAE FullAccess
Buy a package	x	√

Description	CAE ReadOnlyAccess	CAE FullAccess
Create an environment	x	√
Query all environments	√	√
Query environment information	√	√
Delete an environment	x	√
Create an application	x	√
Query application information	√	√
Query all applications	√	√
Update application information	x	√
Delete an application	x	√
Create a component	x	√
Query component information	√	√
Query component configurations	√	√
Query component events	√	√
Query all components and instances	√	√
View usage data	√	√
Modify component configurations	x	√
Scale, upgrade, roll back, stop, start, restart, and edit components	x	√
Delete a component	x	√
Enable certificate configuration	x	√
View certificate configurations	√	√
Modify certificate configurations	x	√
Disable certificate configuration	x	√

Description	CAE ReadOnlyAccess	CAE FullAccess
Configure a domain name	x	√
View a domain name	√	√
Cancel domain name configuration	x	√
Authorize cloud storage	x	√
Unbind cloud storage	x	√
Log in remotely	x	√

Role/Policy Dependencies of the CAE Console

Table 6-3 Role/Policy dependencies of the CAE console

Console Function	Dependency	Role/Policy Required
<ul style="list-style-type: none">Deploy components using OBS software packageAuthorize and unbind cloud storage	Object Storage Service (OBS)	To use these two functions, an IAM user must be granted: CAE FullAccess and OBS Administrator.
Create an environment	Application Operations Management (AOM)	<p>To create an environment, an IAM user must be granted: CAE FullAccess and AOM FullAccess.</p> <p>NOTE</p> <p>If IAM users are granted the V3 system permission CAE FullAccess for an enterprise project, they cannot use the component event and component log functions. Permissions of CAE that depend on AOM do not support enterprise projects, including aom:alarm:list (querying event alarm information), aom:metric:get (querying metrics), and aom:metric:list (querying time series data).</p> <p>For details about AOM permissions policies and supported actions, see Actions Supported by Policy-based Authorization.</p>

Console Function	Dependency	Role/Policy Required
Component creation, deployment, upgrade, and rollback	Software Repository for Container (SWR)	To perform operations on components, an IAM user must be granted: CAE FullAccess and SWR Admin.
Query and configure public domain names	Domain Name Service (DNS)	To query and configure public domain names, an IAM user must be granted: CAE FullAccess and DNS Administrator. NOTE To grant the DNS Administrator permission, the VPC Administrator and Tenant Guest permissions must be granted. After the DNS Administrator permission is granted, the VPC Administrator and Tenant Guest permissions can be deleted. The deletion does not affect domain name configuration.
Buy a package and make payment	Billing Center (BSS)	To buy a package and make payment, an IAM user must be granted: CAE FullAccess and BSS Finance.
<ul style="list-style-type: none">• Configure a secret• Configure environment variables by referring secrets	Data Encryption Workshop (DEW)	To configure secrets and use them to configure environment variables, an IAM user must be granted: CAE FullAccess, KMS CMKFullAccess (all permissions for KMS CMKs), and CSMS ReadonlyAccess (read-only permissions for CSMS).

Customizing Fine-grained Policies

To use a custom fine-grained policy, log in to IAM as the administrator and select fine-grained permissions of CAE as required. [Table 6-4](#) describes fine-grained permission dependencies of CAE.

On the IAM console, choose **Permissions > Policies/Roles > Create Custom Policy > Select service > CAE > Policy Content**. Only the permissions (**cae:environment:*** and **cae:application:***) listed in the following table apply to Huawei Cloud regions.

Table 6-4 Fine-grained permission dependencies of CAE

Permission Name	Description	Permission Dependency	Scenarios (Actions)
cae:environment:create	Create an environment	<ul style="list-style-type: none">vpc:vpcs:getvpc:vpcs:listvpc:vpcs:createvpc:subnets:getvpc:subnets:createvpc:ports:getvpc:ports:createOBS AdministratorAOM FullAccessBSS FinanceSWR Admin	<ul style="list-style-type: none">Create an environmentBuy a packageEnable certificate configurationModify certificate configurationsConfigure a domain nameAuthorize cloud storage
cae:environment:list	Query all environments	<ul style="list-style-type: none">obs:bucket:GetBucketLocationobs:bucket:GetBucketStorageobs:bucket:GetBucketStoragePolicy	<ul style="list-style-type: none">Query an environmentView cloud storage authorization
cae:environment:get	Query environment information	None	<ul style="list-style-type: none">View certificate configurationsView a domain name -
cae:environment:delete	Delete an environment	<ul style="list-style-type: none">vpc:ports:deleteOBS Administrator	<ul style="list-style-type: none">Delete an environmentDisable certificate configurationCancel domain name configurationUnbind cloud storage

Permission Name	Description	Permission Dependency	Scenarios (Actions)
cae:application:create	Create an application	<ul style="list-style-type: none"> rds:instance:list rds:databaseUser:list rds:database:list cse:engine:get cse:engine:list lts:groups:create lts:topics:create lts:*.list lts:*.get OBS Administrator SWR Admin 	<ul style="list-style-type: none"> Create an application Create the components Modify component configurations
cae:application:get	Query application information	<ul style="list-style-type: none"> rds:instance:list rds:databaseUser:list rds:database:list cse:engine:get cse:engine:list 	<ul style="list-style-type: none"> View application information View component information View component configuration
cae:application:list	Query all applications	None	<ul style="list-style-type: none"> View all applications View usage data View all components and instances View component events
cae:application:modify	Update application information	<ul style="list-style-type: none"> OBS Administrator SWR Admin 	Component operations: <ul style="list-style-type: none"> Scaling Upgrade Rollback Stop Start Restart Modify

Permission Name	Description	Permission Dependency	Scenarios (Actions)
cae:application:delete	Delete an application	<ul style="list-style-type: none"> • lts:groups:delete • lts:topics:delete • OBS Administrator 	<ul style="list-style-type: none"> • Delete an application • Delete a component • Delete component configuration
cae:application:createConsole	Log in remotely	<ul style="list-style-type: none"> • cae:application:delete • cae:application:modify • cae:application:get 	Log in remotely

7 Relationship with Other Cloud Services

CAE is a one-stop cloud application platform. Its core service functions cover infrastructure, storage, database, software repository, monitoring and O&M, and middleware services.

CAE:

- Interconnects with source code repositories, such as CodeArts, GitHub, Gitee, GitLab, and Bitbucket. After it is bound, you can directly pull up the source code from source code repositories for building.
- Integrates the software center and archives the built software packages (or image packages) to the corresponding repositories and organizations.
- Integrates Cloud Service Engine (CSE) to register component instances with CSE. CSE provides all-scenario capabilities such as service registration, governance, and configuration management. CSE enables users to quickly develop microservices and perform high-availability O&M.
- Integrates Rational Database Service (RDS) to store data exchanged between application components to RDS. RDS is a professional database management platform for users to easily set up and expand cloud databases.
- Provides application component monitoring, O&M management, and log services.
- Interconnects with Object Storage Service (OBS) for specified functions. CAE enables users to create and query buckets in OBS, and upload, query, and delete object files in buckets. They also can use CAE to create and query parallel file systems in OBS.