AOM

Service Overview

Issue 01

Date 2024-05-10

HUAWEI

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 Is AOM?

Challenges

With the popularization of container technologies, lots of enterprises develop applications using microservice frameworks. Because the number of cloud services increases, enterprises gradually turn to cloud O&M. However, they face the following O&M challenges:

Figure 1-1 Existing O&M issues

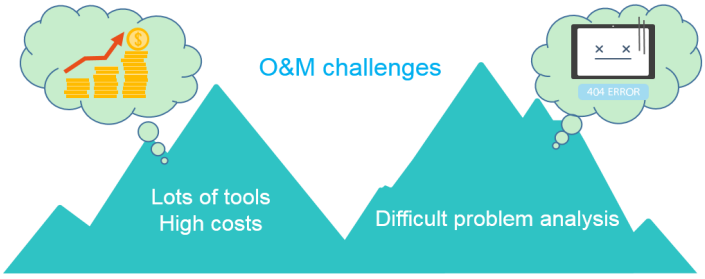

- Cloud O&M has high requirements on personnel skills. O&M tools are hard to configure. Multiple systems need to be maintained at the same time. Distributed tracing systems face high learning and usage costs, but have poor stability.
- Distributed applications face analysis difficulties such as how to visualize the dependency between microservices, improve user experience, associate scattered logs for analysis, and quickly trace problems.

Introduction to AOM

Figure 1-2 One-stop O&M platform

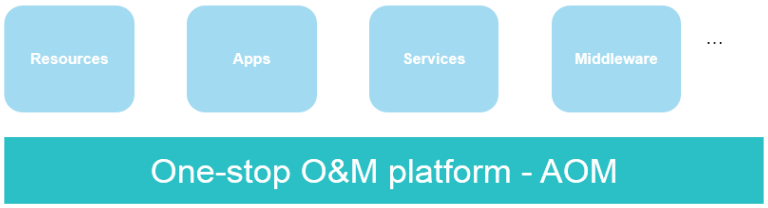

Application Operations Management (AOM) is a one-stop, multi-dimensional O&M management platform for cloud applications. **It monitors your applications** and related cloud resources, analyzes application health status in real time, and provides flexible **data visualization** functions, helping you monitor running status of applications, resources, and services in real time and **detect faults in a timely manner**.

Advantages

Figure 1-3 AOM advantage 1

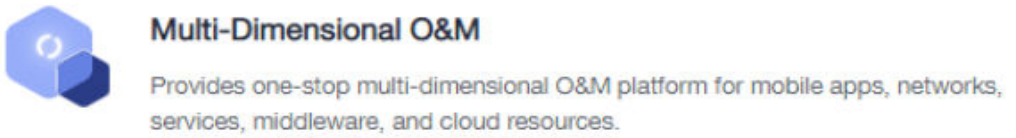

Figure 1-4 AOM advantage 2

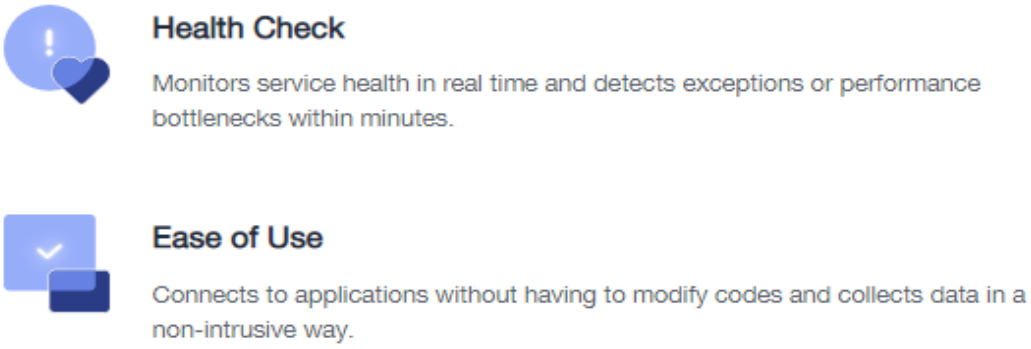

- **Management over massive quantities of logs**

  AOM supports **log search** and service analysis, automatically associates logs for cluster analysis, and filters logs by application, host, file, or instance.
- **Association analysis**

  AOM automatically associates applications and resources and displays data in a panorama view. Through analysis of metrics and alarms about applications, components, instances, hosts, and transactions, AOM allows you to easily locate faults.

- **Open ecosystem**

  O&M data query APIs are opened, collection standards are provided, and independent development is supported.

2 Product Architecture

AOM is a multi-dimensional O&M platform that focuses on resource data and associates log, metric, resource, alarm, and event data. It consists of the data collection and access layer, transmission and storage layer, and service computing layer.

Architecture Diagram

Figure 2-1 AOM architecture

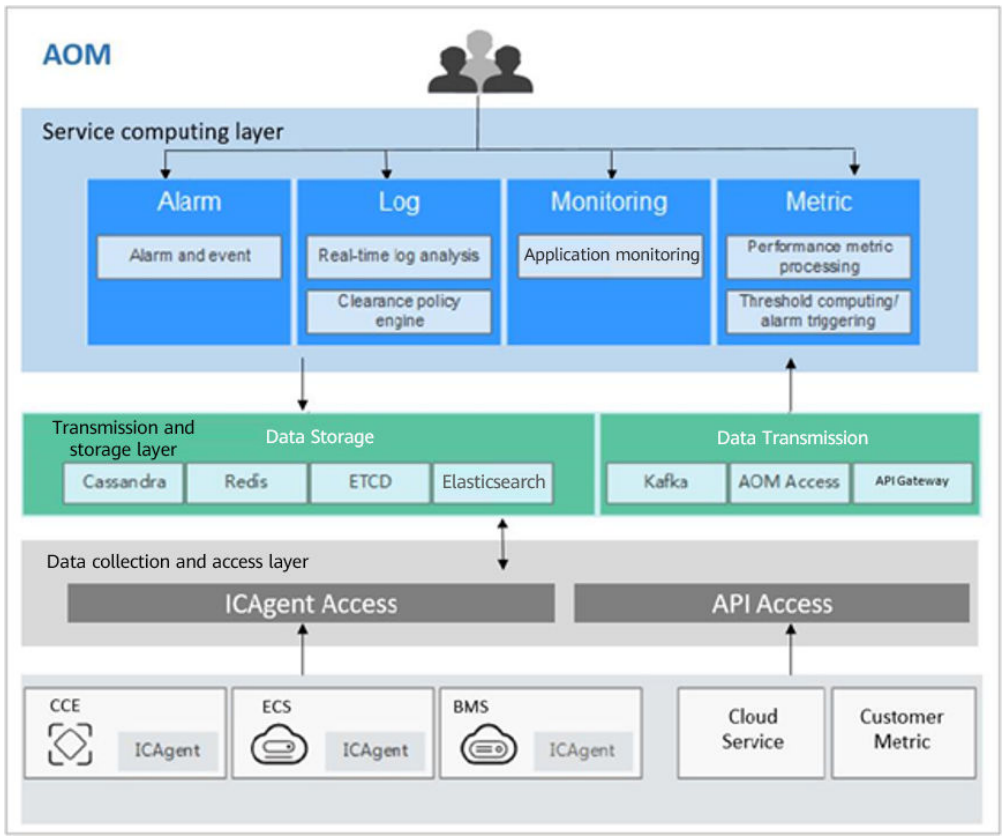

Architecture Description

- **Data collection and access layer**
  - Collecting data by using ICAgent

    You can **install the ICAgent** (a data collector) on a host and use it to report O&M data.
  - Connecting data by using APIs

    You can connect service metrics to AOM as custom metrics using AOM open APIs or Exporter APIs.
- **Transmission and storage layer**
  - Data transmission: AOM Access is a proxy for receiving O&M data. After O&M data is received, such data will be placed in the Kafka queue. Kafka then transmits the data to the service computing layer in real time based on its high-throughput capability.
  - Data storage: After being processed by the AOM backend, O&M data is written into databases. Cassandra stores , Redis is used for cache query, etcd stores AOM configuration data, and Elasticsearch stores resources, logs, alarms, and events.
- **Service computing layer**

  AOM provides basic O&M services such as **alarm management**, **log management**, and resource monitoring (such as **metric monitoring**). It also provides AI services such as exception detection and analysis.

3 Functions

Application Monitoring

Application monitoring allows you to view application resource usage, trends, and alarms in real time, so that you can make fast responses to ensure smooth running for applications.

This function adopts the hierarchical drill-down design. The hierarchy is as follows: Application list > Application details > Component details > Instance details > Process details. Applications, components, instances, and processes are visually associated with each other on the console.

Host Monitoring

Host monitoring allows you to view host resource usage, trends, and alarms in real time, so that you can make fast responses and ensure smooth running for hosts.

Like application monitoring, this function also adopts the hierarchical drill-down design. The hierarchy is as follows: Host list > Host details. The details page contains all the instances, GPUs, NICs, disks, and file systems of the current host.

Automatic Discovery of Applications

After you deploy applications on hosts, the ICAgent installed on the hosts automatically collects information, including names of processes, components, containers, and Kubernetes pods. Applications are **automatically discovered** and their graphs are displayed on the console. You can then set aliases and groups for better resource management.

Dashboard

With a **dashboard**, different graphs can be displayed on the same screen. Various graphs, such as , digit graphs, and top N resource graphs enable you to monitor data comprehensively.

For example, you can add key metrics to a dashboard for real-time monitoring. You can also compare the same metric of different resources on one screen. In addition, by adding common O&M metrics to a dashboard, you do not need to reselect them when re-opening the AOM console during routine O&M.

Alarm Management

The **alarm list** helps you manage alarms and events.

You can **create threshold rules** for key resource metrics. When the metric value reaches the threshold, AOM will generate alarms.

Log Management

AOM provides powerful **log management capabilities**. **Log search** enables you to quickly search for required logs from massive quantities of logs. **Log dump** enables you to store logs for a long time. After you **create log statistical rules**, AOM can periodically count keywords and generate metric data, so that you can monitor system performance and services in real time. By **configuring delimiters**, you can divide log content into multiple words and use these words to search for logs.

Metric Browsing

The **Metric Browsing** page displays metric data of each resource. You can monitor metric values and trends in real time, and create alarm rules for desired metrics. In this way, you can monitor services in real time and perform data correlation analysis.

4 Application Scenarios

AOM is widely used. You can learn how to use AOM in the following typical scenarios.

Problem Inspection and Demarcation

During routine O&M, it is hard to locate faults and obtain logs. Therefore, a monitoring platform is required to monitor resources, logs, and application performance.

AOM interconnects with application services, and collects O&M data of infrastructures, middleware, and application instances in one stop. Through metric monitoring, log analysis, and alarm reporting, AOM enables you to monitor the application running status and resource usage easily, and detect and demarcate problems in a timely manner.

Advantages

- Automatic discovery of applications: Collectors are deployed to proactively discover and monitor applications based on different runtime environments.
- Monitoring of distributed applications: AOM serves as a unified O&M platform that enables you to implement multi-dimensional monitoring over distributed applications with multiple cloud services.
- Alarm notification: Multiple exception detection policies, alarm trigger modes, and APIs are provided.

Figure 4-1 Problem inspection and demarcation

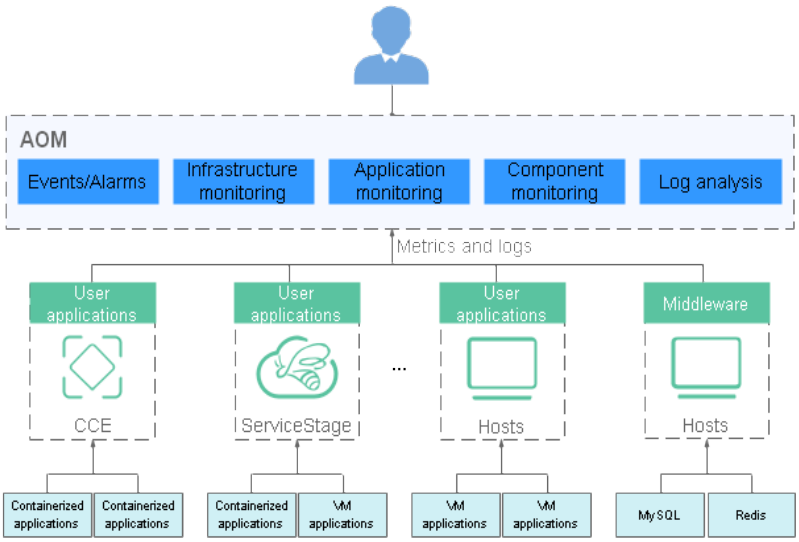

Multi-Dimensional O&M

You need to monitor comprehensive system running status and make fast response to various problems.

AOM provides multi-dimensional O&M capabilities from the cloud level to the resource level and from application monitoring to microservice tracing.

Advantages

- User experience assurance: Service health status KPIs in real time are monitored in real time and root causes of exceptions are analyzed.
- Fast fault diagnosis: Distributed call tracing enables you to locate faults quickly.
- Resource running assurance: Hundreds of O&M metrics about resources such as containers, disks, and networks are monitored in real time, and clusters, VMs, applications, and containers are associated for analysis.

Figure 4-2 Multi-dimensional O&M

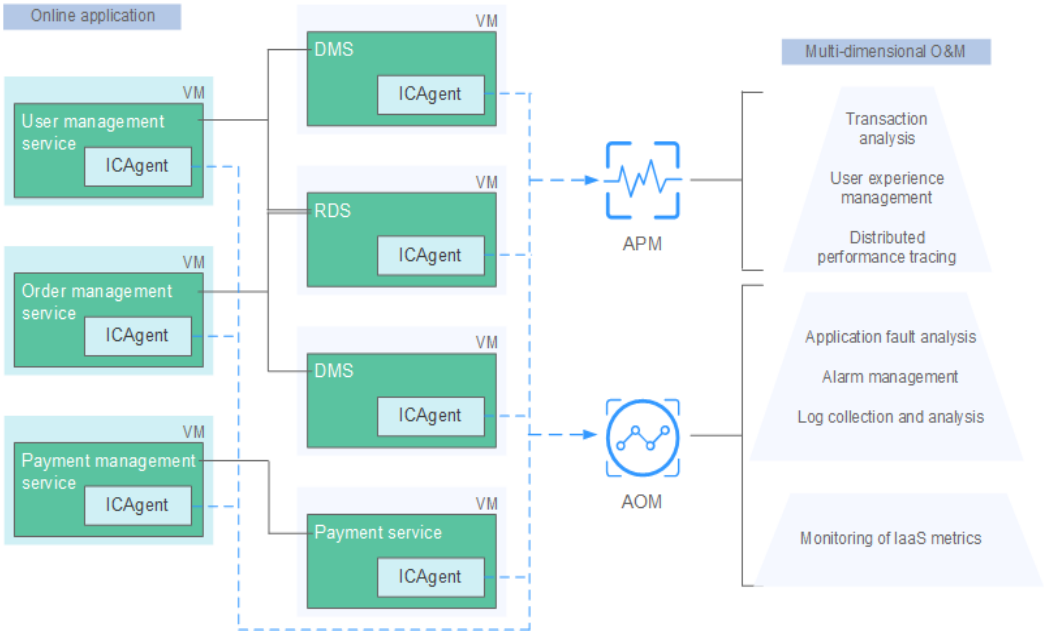

5 Edition Differences

AOM provides the basic and pay-per-use editions. The pay-per-use edition can be further classified into the professional edition and enterprise edition.

For details about the edition differences, see **Table 5-1**.

Table 5-1 Edition differences

Item	Basic Edition	Professional Edition	Enterprise Edition
Log read and write traffic	500 MB	400 GB	1 TB
Log index traffic	500 MB	400 GB	1 TB
Log storage space	500 MB	400 GB	1 TB
Host monitoring	Metric storage duration: 7 days	100 VMs; metric storage duration: 1 year	200 VMs; metric storage duration: 1 year
CCI instance monitoring	Metric storage duration: 7 days	500 instances; metric storage duration: 1 year	1000 instances; metric storage duration: 1 year
Metrics in real-time monitoring	Free of charge in the first month: 10	N/A	2000
Custom metrics	10	100	500
API call (metrics query)	500,000 calls/ month	1 million calls/ month	5 million calls/ month
Events and alarms	50,000 pieces/ month; storage period: 7 days	100,000 pieces/ month; storage period: 30 days	500,000 pieces/ month; storage period: 30 days

6 Metric Overview

6.1 Introduction

Metrics reflect resource performance data or status. A metric consists of a **namespace**, **dimension**, name, and unit. Metrics can be divided into:

- System metrics: basic metrics provided by AOM, such as CPU usage and used CPU cores.
- Custom metrics: user-defined metrics. Custom metrics can be reported using the following methods:
  - Method 1: Use AOM APIs. For details, see **Adding Monitoring Data** and **Querying Monitoring Data**.
  - Method 2: When creating containerized applications on CCE, interconnect with Prometheus to report custom metrics. For details, see **Custom Monitoring**.

Metric Namespaces

A namespace is an abstract collection of resources and objects. Metrics in different namespaces are independent of each other so that metrics of different applications will not be aggregated to the same statistics information.

- Namespaces of system metrics are fixed and started with **PAAS.**. For details, see **Table 6-1**.

Table 6-1 Namespaces of system metrics

Namespace	Description
PAAS.AGGR	Namespace of cluster metrics
PAAS.NODE	Namespace of host, network, disk, and file system metrics
PAAS.CONTAINER	Namespace of component, instance, process, and container metrics
PAAS.SLA	Namespace of SLA metrics

- Namespaces of custom metrics must be in the XX.XX format. Each namespace must be 3 to 32 characters long, starting with a letter (excluding **PAAS.**, **SYS.**, and **SRE.**). Only digits, letters, and underscores (_) are allowed.

Metric Dimensions

Metric dimensions indicate the categories of metrics. Each metric has certain features, and a dimension may be considered as a category of such features.

- Dimensions of system metrics are fixed. Different types of metrics have different dimensions. For more details, see the following sections.
- Dimensions of custom metrics must be 1 to 32 characters long, which need to be customized.

6.2 Network Metrics and Dimensions

Table 6-2 Network metrics

Metric	Description	Value Range	Unit
Downlink rate (BPS) (aom_node_network_receive_bytes)	Inbound traffic rate of a measured object	≥ 0	Byte/s
Downlink rate (PPS) (aom_node_network_receive_packets)	Number of data packets received by a NIC per second	≥ 0	Packet/s
Downlink error rate (aom_node_network_receive_error_packets)	Number of error packets received by a NIC per second	≥ 0	Count/s
Uplink rate (BPS) (aom_node_network_transmit_bytes)	Outbound traffic rate of a measured object	≥ 0	Byte/s
Uplink error rate (aom_node_network_transmit_error_packets)	Number of error packets sent by a NIC per second	≥ 0	Count/s
Uplink rate (PPS) (aom_node_network_transmit_packets)	Number of data packets sent by a NIC per second	≥ 0	Packet/s
Total rate (BPS) (aom_node_network_total_bytes)	Total inbound and outbound traffic rate of a measured object	≥ 0	Byte/s

Table 6-3 Dimensions of network metrics

Dimension	Description
clusterId	Cluster ID
hostID	Host ID
nameSpace	Cluster namespace
netDevice	NIC name
nodeIP	Host IP address
nodeName	Host name

6.3 Disk Metrics and Dimensions

Table 6-4 Disk metrics

Metric	Description	Value Range	Unit
Disk read rate (aom_node_disk_read_kilobytes)	Volume of data read from a disk per second	≥ 0	KB/s
Disk write rate (aom_node_disk_write_kilobytes)	Volume of data written into a disk per second	≥ 0	KB/s

Table 6-5 Dimensions of disk metrics

Dimension	Description
clusterId	Cluster ID
diskDevice	Disk name
hostID	Host ID
nameSpace	Cluster namespace
nodeIP	Host IP address
nodeName	Host name

6.4 Disk Partition Metrics

NOTE

- If the host type is **CCE**, you can view disk partition metrics. The supported OSs are CentOS 7.6 and EulerOS 2.5.
- Log in to the CCE node as the **root** user and run the **docker info | grep 'Storage Driver'** command to check the Docker storage driver type. If the command output shows driver type **Device Mapper**, the thin pool metrics can be viewed. Otherwise, the thin pool metrics cannot be viewed.

Table 6-6 Disk partition metrics

Metric	Description	Value Range	Unit
Thin pool's metadata space usage (aom_host_diskpartition_thinpool_metadata_percent)	Percentage of the thin pool's used metadata space to the total metadata space on a CCE node	0–100	%
Thin pool's data space usage (aom_host_diskpartition_thinpool_data_percent)	Percentage of the thin pool's used data space to the total data space on a CCE node	0–100	%
Thin pool's disk partition space (aom_host_diskpartition_total_capacity_megabytes)	Total thin pool's disk partition space on a CCE node	≥ 0	MB

6.5 File System Metrics and Dimensions

Table 6-7 File system metrics

Metric	Description	Value Range	Unit
Available disk space (aom_node_disk_available_capacity_megabytes)	Disk space that has not been used	≥ 0	MB
Total disk space (aom_node_disk_capacity_megabytes)	Total disk space	≥ 0	MB

Metric	Description	Value Range	Unit
Disk read/write status (aom_node_disk_rw_status)	Read or write status of a disk	0 or 1 • 0: read/write • 1: read-only	N/A
Disk usage (aom_node_disk_usage)	Percentage of the used disk space to the total disk space	0–100	%

Table 6-8 Dimensions of file system metrics

Dimension	Description
clusterId	Cluster ID
clusterName	Cluster name
fileSystem	File system
hostID	Host ID
mountPoint	Mount point
nameSpace	Cluster namespace
nodeIP	Host IP address
nodeName	Host name

6.6 Host Metrics and Dimensions

Table 6-9 Host metrics

Metric	Description	Value Range	Unit
Total CPU cores (aom_node_cpu_limit_core)	Total number of CPU cores that have been applied for a measured object	≥ 1	Cores

Metric	Description	Value Range	Unit
Used CPU cores (aom_node_cpu_used_core)	Number of CPU cores used by a measured object	≥ 0	Cores
CPU usage (aom_node_cpu_usage)	CPU usage of a measured object	0–100	%
Available physical memory (aom_node_memory_free_megabytes)	Available physical memory of a measured object	≥ 0	MB
Available virtual memory (aom_node_virtual_memory_free_megabytes)	Available virtual memory of a measured object	≥ 0	MB
Total GPU memory (aom_node_gpu_memory_free_megabytes)	Total GPU memory of a measured object	> 0	MB
GPU memory usage (aom_node_gpu_memory_usage)	Percentage of the used GPU memory to the total GPU memory	0–100	%
Used GPU memory (aom_node_gpu_memory_used_megabytes)	GPU memory used by a measured object	≥ 0	MB
GPU usage (aom_node_gpu_usage)	GPU usage of a measured object	0–100	%
Total NPU memory (aom_node_npu_memory_free_megabytes)	Total NPU memory of a measured object	> 0	MB
NPU memory usage (aom_node_npu_memory_usage)	Percentage of the used NPU memory to the total NPU memory	0–100	%
Used NPU memory (aom_node_npu_memory_used_megabytes)	NPU memory used by a measured object	≥ 0	MB
NPU usage (aom_node_npu_usage)	NPU usage of a measured object	0–100	%
NPU temperature (aom_node_npu_temperature_centigrade)	NPU temperature of a measured object	-	°C

Metric	Description	Value Range	Unit
Physical memory usage (aom_node_memory_usage)	Percentage of the used physical memory to the total physical memory	0–100	%
Host status (aom_node_status)	Host status	• 0: Normal • 1: Abnormal	N/A
NTP offset (aom_node_ntp_offset_ms)	Offset between the local time of the host and the NTP server time. The closer the NTP offset is to 0, the closer the local time of the host is to the time of the NTP server.	-	ms
NTP server status (aom_node_ntp_server_status)	Whether the host is connected to the NTP server	0 or 1 • 0: Connected • 1: Unconnected	N/A
NTP synchronization status (aom_node_ntp_status)	Whether the local time of the host is synchronized with the NTP server time	0 or 1 • 0: Synchronous • 1: Not synchronized	N/A
Processes (aom_node_process_number)	Number of processes on a measured object	≥ 0	N/A
GPU temperature (aom_node_gpu_temperature_centigrade)	GPU temperature of a measured object	-	°C

Metric	Description	Value Range	Unit
Total physical memory (aom_node_memory_total_megabytes)	Total physical memory that has been applied for a measured object	≥ 0	MB
Total virtual memory (aom_node_virtual_memory_total_megabytes)	Total virtual memory that has been applied for a measured object	≥ 0	MB
Virtual memory usage (aom_node_virtual_memory_usage)	Percentage of the used virtual memory to the total virtual memory	0–100	%
Threads (aom_node_current_threads_num)	Number of threads created on a host	≥ 0	N/A
Max. threads (aom_node_sys_max_threads_num)	Maximum number of threads that can be created on a host	≥ 0	N/A
Total physical disk space (aom_node_phy_disk_total_capacity_megabytes)	Total disk space of a host	≥ 0	MB
Used disk space (aom_node_physical_disk_total_used_megabytes)	Used disk space of a host	≥ 0	MB
Hosts (aom_billing_hostUsed)	Number of hosts connected per day	≥ 0	N/A

NOTE

- Memory usage = (Physical memory capacity – Available physical memory capacity)/Physical memory capacity; Virtual memory usage = ((Physical memory capacity + Total virtual memory capacity) – (Available physical memory capacity + Available virtual memory capacity))/(Physical memory capacity + Total virtual memory capacity)
- The virtual memory of a VM is 0 MB by default. If no virtual memory is configured, the memory usage on the monitoring page is the same as the virtual memory usage.
- For the total and used physical disk space, only the space of the local disk partitions' file systems is counted. The file systems (such as JuiceFS, NFS, and SMB) mounted to the host through the network are not taken into account.

Table 6-10 Dimensions of host metrics

Dimension	Description
clusterId	Cluster ID
clusterName	Cluster name
gpuName	GPU name
gpuID	GPU ID
npuName	NPU name
npuID	NPU ID
hostID	Host ID
nameSpace	Cluster namespace
nodeIP	Host IP address
hostName	Host name

6.7 Cluster Metrics and Dimensions

NOTE

Cluster metrics are aggregated by AOM based on host metrics, and do not include the metrics of master nodes.

Table 6-11 Cluster metrics

Metric	Description	Value Range	Unit
Total CPU cores (aom_cluster_cpu_limit_core)	Total number of CPU cores that have been applied for a measured object	≥ 1	Cores
Used CPU cores (aom_cluster_cpu_used_core)	Number of CPU cores used by a measured object	≥ 0	Cores
CPU usage (aom_cluster_cpu_usage)	CPU usage of a measured object	0–100	%
Available disk space (aom_cluster_disk_available_capacity_megabytes)	Disk space that has not been used	≥ 0	MB

Metric	Description	Value Range	Unit
Total disk space (aom_cluster_disk_capacity_megabytes)	Total disk space	≥ 0	MB
Disk usage (aom_cluster_disk_usage)	Percentage of the used disk space to the total disk space	0–100	%
Available physical memory (aom_cluster_memory_free_megabytes)	Available physical memory of a measured object	≥ 0	MB
Available virtual memory (aom_cluster_virtual_memory_free_megabytes)	Available virtual memory of a measured object	≥ 0	MB
Available GPU memory (aom_cluster_gpu_memory_free_megabytes)	Available GPU memory of a measured object	> 0	MB
GPU memory usage (aom_cluster_gpu_memory_usage)	Percentage of the used GPU memory to the total GPU memory	0–100	%
Used GPU memory (aom_cluster_gpu_memory_used_megabytes)	GPU memory used by a measured object	≥ 0	MB
GPU usage (aom_cluster_gpu_usage)	GPU usage of a measured object	0–100	%
Physical memory usage (aom_cluster_memory_usage)	Percentage of the used physical memory to the total physical memory	0–100	%
Downlink rate (BPS) (aom_cluster_network_receive_bytes)	Inbound traffic rate of a measured object	≥ 0	Byte/s
Uplink rate (BPS) (aom_cluster_network_transmit_bytes)	Outbound traffic rate of a measured object	≥ 0	Byte/s
Total physical memory (aom_cluster_memory_total_megabytes)	Total physical memory that has been applied for a measured object	≥ 0	MB

Metric	Description	Value Range	Unit
Total virtual memory (aom_cluster_virtual_memory_total_megabytes)	Total virtual memory that has been applied for a measured object	≥ 0	MB
Virtual memory usage (aom_cluster_virtual_memory_usage)	Percentage of the used virtual memory to the total virtual memory	0–100	%

Table 6-12 Dimensions of cluster metrics

Dimension	Description
clusterId	Cluster ID
clusterName	Cluster name
projectId	Project ID

6.8 Container Metrics and Dimensions

Table 6-13 Container metrics

Metric	Description	Value Range	Unit
Total CPU cores (aom_container_cpu_limit_core)	Total number of CPU cores restricted for a measured object	≥ 1	Cores
Used CPU cores (aom_container_cpu_used_core)	Number of CPU cores used by a measured object	≥ 0	Cores
CPU usage (aom_container_cpu_usage)	CPU usage of a measured object. That is, the percentage of the used CPU cores to the total CPU cores restricted for a measured object.	0–100	%
Disk read rate (aom_container_disk_read_kilobytes)	Volume of data read from a disk per second	≥ 0	KB/s
Disk write rate (aom_container_disk_write_kilobytes)	Volume of data written into a disk per second	≥ 0	KB/s

Metric	Description	Value Range	Unit
Available file system capacity (aom_container_filesystem_available_capacity_megabytes)	Available file system capacity of a measured object. This metric is available only for containers using the Device Mapper storage drive in the Kubernetes cluster of version 1.11 or later.	≥ 0	MB
Total file system capability (aom_container_filesystem_capacity_megabytes)	Total file system capacity of a measured object. This metric is available only for containers using the Device Mapper storage drive in the Kubernetes cluster of version 1.11 or later.	≥ 0	MB
File system usage (aom_container_filesystem_usage)	File system usage of a measured object. That is, the percentage of the used file system to the total file system. This metric is available only for containers using the Device Mapper storage drive in the Kubernetes cluster of version 1.11 or later.	0–100	%
Total GPU memory (aom_container_gpu_memory_free_megabytes)	Total GPU memory of a measured object	> 0	MB
GPU memory usage (aom_container_gpu_memory_usage)	Percentage of the used GPU memory to the total GPU memory	0–100	%
Used GPU memory (aom_container_gpu_memory_used_megabytes)	GPU memory used by a measured object	≥ 0	MB
GPU usage (aom_container_gpu_usage)	GPU usage of a measured object	0–100	%
Total NPU memory (aom_container_npu_memory_free_megabytes)	Total NPU memory of a measured object	> 0	MB
NPU memory usage (aom_container_npu_memory_usage)	Percentage of the used NPU memory to the total NPU memory	0–100	%

Metric	Description	Value Range	Unit
Used NPU memory (aom_container_npu_memory_used_megabytes)	NPU memory used by a measured object	≥ 0	MB
NPU usage (aom_container_npu_usage)	NPU usage of a measured object	0–100	%
Total physical memory (aom_container_memory_request_megabytes)	Total physical memory restricted for a measured object	≥ 0	MB
Physical memory usage (aom_container_memory_usage)	Percentage of the used physical memory to the total physical memory restricted for a measured object	0–100	%
Used physical memory (aom_container_memory_used_megabytes)	Used physical memory of a measured object	≥ 0	MB
Downlink rate (BPS) (aom_container_network_receive_bytes)	Inbound traffic rate of a measured object	≥ 0	Byte/s
Downlink rate (PPS) (aom_container_network_receive_packets)	Number of data packets received by a NIC per second	≥ 0	Packet/s
Downlink error rate (aom_container_network_receive_error_packets)	Number of error packets received by a NIC per second	≥ 0	Count/s
Error packets (aom_container_network_rx_error_packets)	Number of error packets received by a measured object	≥ 0	Count
Uplink rate (BPS) (aom_container_network_transmit_bytes)	Outbound traffic rate of a measured object	≥ 0	Byte/s
Uplink error rate (aom_container_network_transmit_error_packets)	Number of error packets sent by a NIC per second	≥ 0	Count/s
Uplink rate (PPS) (aom_container_network_transmit_packets)	Number of data packets sent by a NIC per second	≥ 0	Packet/s

Metric	Description	Value Range	Unit
Status (aom_process_status)	Docker container status	0 or 1 • 0: Normal • 1: Abnormal	N/A
Working set memory usage (aom_container_memory_workingset_usage)	Usage of the working set memory	0–100	%
Used working set memory (aom_container_memory_workingset_used_megabytes)	Sum of resident set size (RSS) memory and cache	≥ 0	MB

Table 6-14 Dimensions of container metrics

Dimension	Description
appID	Service ID
appName	Service name
clusterId	Cluster ID
clusterName	Cluster name
containerID	Container ID
containerName	Container name
deploymentName	Kubernetes deployment name
kind	Application type
nameSpace	Cluster namespace
podID	Instance ID
podName	Instance name
serviceID	Inventory ID
gpuID	GPU ID
npuName	NPU name

Dimension	Description
npuID	NPU ID

6.9 VM Metrics and Dimensions

In AOM, VMs refer to processes, and VM metrics refer to process metrics.

Table 6-15 Process metrics

Metric	Description	Value Range	Unit
Total CPU cores (aom_process_cpu_limit_core)	Total number of CPU cores that have been applied for a measured object	≥ 1	Cores
Used CPU cores (aom_process_cpu_used_core)	Number of CPU cores used by a measured object	≥ 0	Cores
CPU usage (aom_process_cpu_usage)	CPU usage of a measured object. That is, the percentage of the used CPU cores to the total CPU cores.	0–100	%
Handles (aom_process_handle_count)	Number of handles used by a measured object	≥ 0	N/A
Max. handles (aom_process_max_handle_count)	Maximum number of handles used by a measured object	≥ 0	N/A
Total physical memory (aom_process_memory_request_megabytes)	Total physical memory that has been applied for a measured object	≥ 0	MB
Physical memory usage (aom_process_memory_usage)	Percentage of the used physical memory to the total physical memory	0–100	%
Used physical memory (aom_process_memory_used_megabytes)	Used physical memory of a measured object	≥ 0	MB

Metric	Description	Value Range	Unit
Status (aom_process_status)	Process status	0 or 1 • 0: Normal • 1: Abnormal	N/A
Threads (aom_process_thread_count)	Number of threads used by a measured object	≥ 0	N/A
Total virtual memory (aom_process_virtual_memory_total_megabytes)	Total virtual memory that has been applied for a measured object	≥ 0	MB

Table 6-16 Dimensions of process metrics

Dimension	Description
appName	Service name
clusterId	Cluster ID
clusterName	Cluster name
nameSpace	Cluster namespace
processID	Process ID
processName	Process name
serviceID	Inventory ID
aomApplicationName	Application name
aomApplicationID	Application ID
processCmd	Process command ID

6.10 Instance Metrics and Dimensions

Instance metrics consist of container or process metrics. The dimensions of instance metrics are the same as those of container or process metrics. For details, see **Container Metrics and Dimensions** and **VM Metrics and Dimensions**.

6.11 Service Metrics and Dimensions

Service metrics consist of instance metrics. The dimensions of service metrics are the same as those of instance metrics. For details, see **Instance Metrics and Dimensions**.

7 Security

7.1 Shared Responsibilities

Huawei guarantees that its commitment to cyber security will never be outweighed by the consideration of commercial interests. To cope with emerging cloud security challenges and pervasive cloud security threats and attacks, Huawei Cloud builds a comprehensive cloud service security assurance system for different regions and industries based on Huawei's unique software and hardware advantages, laws, regulations, industry standards, and security ecosystem.

Figure 7-1 illustrates the responsibilities shared by Huawei Cloud and users.

- **Huawei Cloud**: Ensure the security of cloud services and provide secure clouds. Huawei Cloud's security responsibilities include ensuring the security of our IaaS, PaaS, and SaaS services, as well as the physical environments of the Huawei Cloud data centers where our IaaS, PaaS, and SaaS services operate. Huawei Cloud is responsible for not only the security functions and performance of our infrastructure, cloud services, and technologies, but also for the overall cloud O&M security and, in the broader sense, the security and compliance of our infrastructure and services.
- **Tenant**: Use the cloud securely. Tenants of Huawei Cloud are responsible for the secure and effective management of the tenant-customized configurations of cloud services including IaaS, PaaS, and SaaS. This includes but is not limited to virtual networks, the OS of virtual machine hosts and guests, virtual firewalls, API Gateway, advanced security services, all types of cloud services, tenant data, identity accounts, and key management.

Huawei Cloud Security White Paper elaborates on the ideas and measures for building Huawei Cloud security, including cloud security strategies, the shared responsibility model, compliance and privacy, security organizations and personnel, infrastructure security, tenant service and security, engineering security, O&M security, and ecosystem security.

Figure 7-1 Huawei Cloud shared security responsibility model

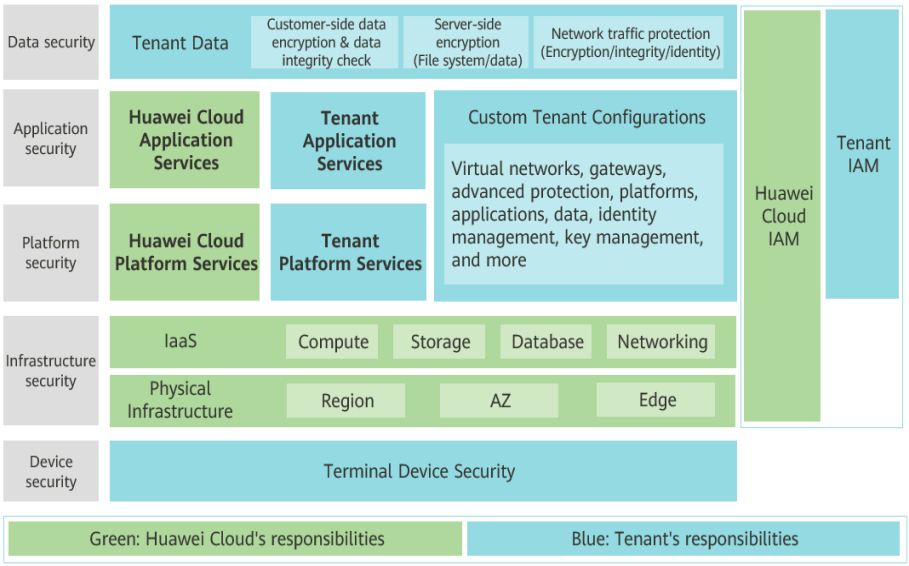

7.2 Identity Authentication and Access Control

Identity Authentication

You are required to present your identity credential and undergo identity authentication no matter whether you access AOM through the console or by calling APIs. In addition, login protection and login authentication policies are provided to harden identity authentication security. Based on Identity and Access Management (IAM), AOM supports identity authentication using **passwords**, **access keys**, or **temporary access keys (for federated users)**. **Login protection** and **login authentication policies** are also provided.

Access Control

If you need to assign different permissions to employees in your enterprise to access your AOM resources, IAM is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you secure access to your cloud resources. For details, see **Permissions Management**.

7.3 Data Protection

AOM takes different measures to keep data secure and reliable.

Table 7-1 AOM data protection methods and features

Measure	Description	Reference
Transmission encryption (HTTPS)	AOM supports HTTPS to enhance data transmission security.	**Making an API Request**
Data redundancy	Metric, alarm, and configuration data is stored in multiple copies to ensure data reliability.	/
Data subscription	With data subscription enabled, AOM sends your metric and alarm data to specified DMS instances. You can then process the dumped data as you want.	**Data Subscription**

7.4 Audit and Logs

Audit

Cloud Trace Service (CTS) records operations on the cloud resources in your account. You can use the logs generated by CTS to perform security analysis, trace resource changes, audit compliance, and locate faults.

After you enable CTS and configure a tracker, CTS records management traces of AOM for auditing.

For details about how to enable and configure CTS, see **Enabling CTS**.

For the management traces of AOM that can be recorded by CTS, see **Operations Logged by CTS**.

Figure 7-2 CTS

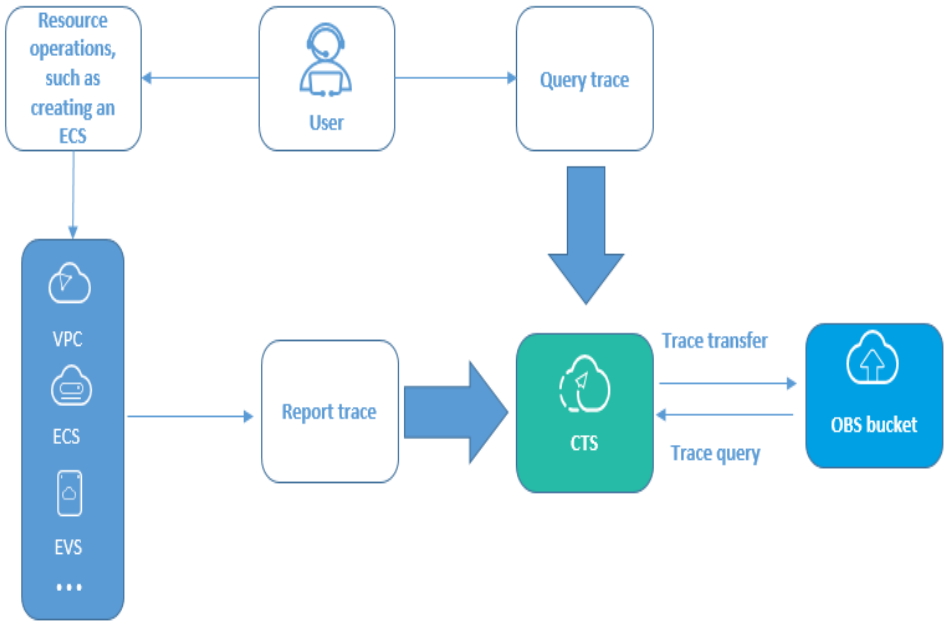

Logs

AOM collects container service logs and VM (ECS or BMS running Linux) logs and displays them on the console for you to search and view. For details, see **Log Management Description**.

7.5 Resilience

AOM provides multiple reliability DR capabilities. Technical solutions (such as intra-AZ instance DR, cross-AZ DR, cross-cluster DR, and multiple data copies) ensure service durability and reliability.

Table 7-2 Reliability architecture of AOM

Reliability Solution	Description
Intra-AZ instance DR	In a single AZ, multiple instances are used for DR. Faulty nodes can be quickly detected and remaining instances can still provide services.
Multi-AZ DR	AOM supports cross-AZ DR. When an AZ is abnormal, instances in other AZs can still provide services.
Cross-cluster DR	AOM supports cross-cluster DR. When one cluster is abnormal, AOM can continue to provide services.
Data DR	AOM configuration, metric, and alarm data is stored in multiple copies to ensure data reliability.

7.6 Security Risk Monitoring

AOM monitors security risks in various ways to ensure data security and reliability. For details, see **Table 7-3**.

Table 7-3 Monitoring security risks

Security Risk Monitoring	Description	Reference
Resource monitoring	AOM supports application, component, host, container, and metric monitoring. It monitors your applications and cloud resources in real time and displays data in a visualized manner, helping you quickly analyze application health status.	**Resource Monitoring Description**
Alarm management	AOM allows you to set alarm conditions for applications, resources, and services based on alarm rules. When AOM or its external service is abnormal or may be abnormal, email, SMS, or WeCom notifications will be sent to specified personnel.	**Alarm Management**

7.7 Certificates

Compliance Certificates

Huawei Cloud services and platforms have obtained various security and compliance certifications from authoritative organizations, such as International Organization for Standardization (ISO). You can **download** them from the console.

Figure 7-3 Downloading compliance certificates

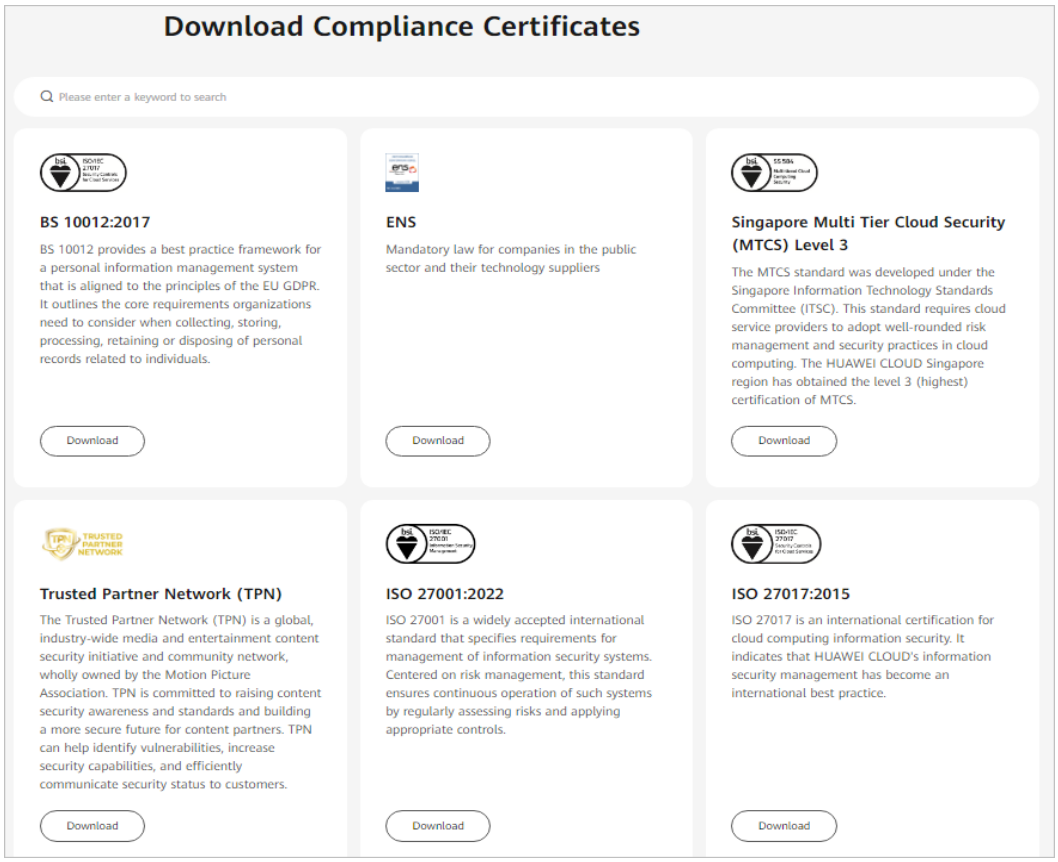

Resource Center

Huawei Cloud also provides the following resources to help users meet compliance requirements. For details, see **Resource Center**.

Figure 7-4 Resource center

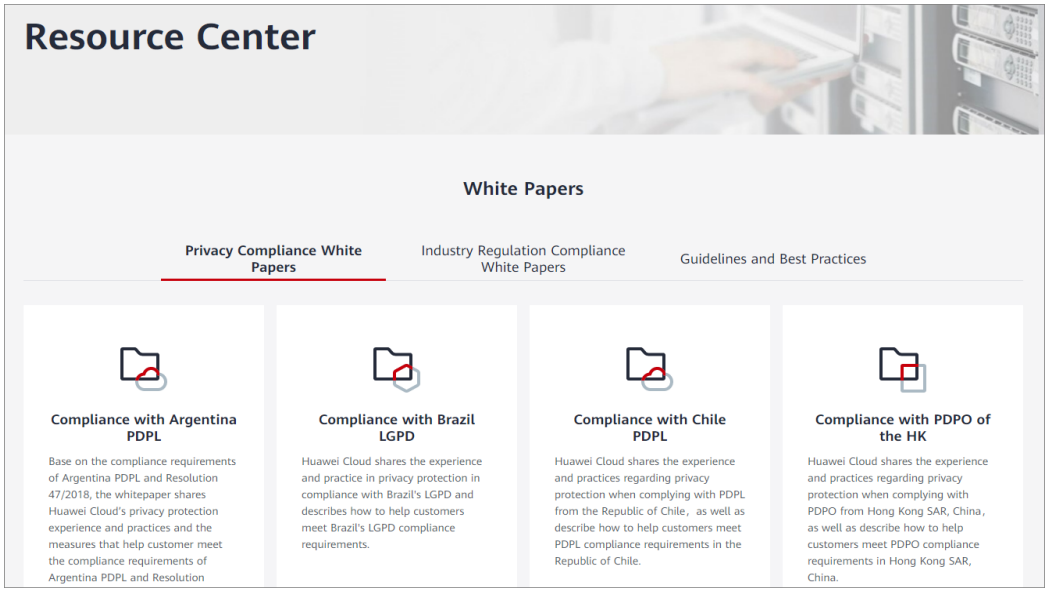

8 Restrictions

OS Usage Restrictions

AOM supports multiple operating systems (OSs). When purchasing a host, ensure that its OS meets the requirements in **Table 8-1**. Otherwise, the host cannot be monitored by AOM.

Table 8-1 OSs and versions supported by AOM

OS	Version					
SUSE	SUSE Enterprise 11 SP4 64-bit	SUSE Enterprise 12 SP1 64-bit	SUSE Enterprise 12 SP2 64-bit	SUSE Enterprise 12 SP3 64-bit		
openSUSE	13.2 64-bit	42.2 64-bit	15.0 64-bit (Currently, syslog logs cannot be collected.)			
EulerOS	2.2 64-bit	2.3 64-bit	2.5 64-bit	2.9 64-bit	2.10 64-bit	
CentOS	6.3 64-bit	6.5 64-bit	6.8 64-bit	6.9 64-bit	6.10 64-bit	
	7.1 64-bit	7.2 64-bit	7.3 64-bit	7.4 64-bit	7.5 64-bit	7.6 64-bit
Ubuntu	14.04 server 64-bit	16.04 server 64-bit	18.04 server 64-bit			
Fedora	24 64-bit	25 64-bit	29 64-bit			
Debian	7.5.0 32-bit	7.5.0 64-bit	8.2.0 64-bit	8.8.0 64-bit	9.0.0 64-bit	
Kylin	Kylin V10 SP1 64-bit					

NOTE

- For Linux x86_64 servers, AOM supports all the OSs and versions listed in the preceding table.
- For Linux Arm servers, AOM only supports CentOS 7.4 and later versions, and other OSs and versions listed in the preceding table.

Resource Usage Restrictions

When using AOM, learn about the restrictions in **Table 8-2**. Resource usage restrictions include quota restrictions. For details, see **Quotas**.

Table 8-2 Resource usage restrictions

Category	Object	Usage Restrictions
Dashboard	Dashboard	A maximum of 50 dashboards can be created in a region.
	Graph in a dashboard	A maximum of 20 graphs can be added to a dashboard.
	Number of resources, threshold rules, components, or hosts in a graph	• A maximum of 100 resources across clusters can be added to a line graph. • A maximum of 12 resources can be added to a digit graph. Only one resource can be displayed. By default, the first resource is displayed. • A maximum of 10 threshold rules can be added to a threshold status graph. • A maximum of 10 hosts can be added to a host status graph. • A maximum of 10 components can be added to a component status graph.
Metric	Metric data	• Basic edition: Metric data can be stored in the database for a maximum of 7 days. • Professional edition: Metric data can be stored in the database for a maximum of 30 days.
	Total number of metrics	Up to 400,000 for a single account. Up to 100,000 for a small specification.
	Metric item	After resources such as clusters, components, and hosts are deleted, their related metrics can be stored in the database for a maximum of 30 days.
	Dimension	A maximum of 20 dimensions can be configured for a metric.

Category	Object	Usage Restrictions
	Metric query API	A maximum of 20 metrics can be queried at a time.
	Statistical period	The maximum statistical period is 1 hour.
	Data points returned for a single query	A maximum of 1440 data points can be returned each time.
	Custom metric	Unlimited.
	Custom metric to be reported	A single request cannot exceed 40 KB. The timestamp of a reported metric cannot 10 minutes later than the standard UTC time. In addition, out-of-order metrics are not received. That is, if a metric is reported at a certain time point, the metrics of earlier time points cannot be reported.
	Application metric	• When the number of containers on a host exceeds 1000, the ICAgent stops collecting application metrics and sends the **ICAgent Stopped Collecting Application Metrics** alarm (ID: 34105). • When the number of containers on a host is less than 1000, the ICAgent resumes the collection of application metrics and the **ICAgent Stopped Collecting Application Metrics** alarm is cleared.
	Resources consumed by the ICAgent	When the ICAgent collects basic metrics, the resources consumed by the ICAgent are greatly affected by the number of containers and processes. On a VM without any services, the ICAgent consumes 30 MB memory and records 1% CPU usage. To ensure collection reliability, ensure that the number of containers running on a single node must be less than 1000.
Threshold rule (Available in AF-Johannesburg, LA-Mexico City1, LA-Mexico City2, LA-Sao Paulo1, and LA-Santiago)	Threshold rule	A maximum of 1000 threshold rules can be created in a project.
	Number of topics that can be selected	A maximum of five topics can be selected for each threshold rule.

Category	Object	Usage Restrictions
Alarm rule (Available in CN North-Beijing1, CN North-Beijing4, CN East-Shanghai1, CN East-Shanghai2, CN South-Guangzhou, CN Southwest-Guiyang1, CN-Hong Kong, AP-Bangkok, and AP-Singapore)	Alarm rule	A maximum of 1000 alarm rules (including threshold rules and event alarm rules) can be created.
	Static threshold template	A maximum of 50 static threshold templates can be created.
Notification rule (available in AF-Johannesburg, LA-Mexico City1, LA-Mexico City2, LA-Sao Paulo1, and LA-Santiago)	Number of topics that can be selected	A maximum of five topics can be selected for each notification rule.
Log	Size of a log	The maximum size of each log is 10 KB. If a log is greater than that, the ICAgent will not collect it. In that case, the log will be discarded.
	Log traffic	A maximum of 10 MB/s is supported for each tenant in a region. If the log traffic exceeds 10 MB/s, logs may be lost. If you require more log traffic, submit a service ticket by following the instructions provided in **Creating a Service Ticket**.
	Log file	Only text log files can be collected. Other types of log files, such as binary files, cannot be collected.
		The ICAgent can collect a maximum of 20 log files from a volume mounting directory.
		The ICAgent can collect a maximum of 1000 standard container output log files. These files must be in JSON format.

Category	Object	Usage Restrictions
	Resources consumed during log file collection	The resources consumed during log file collection are closely related to the log volume, number of files, network bandwidth, and backend service processing capability.
	Log loss	ICAgent uses multiple mechanisms to ensure log collection reliability and prevent data loss. However, logs may be lost in the following scenarios: • The log rotation policy of Cloud Container Engine (CCE) is not used. • Log files are rotated at a high speed, for example, once per second. • Logs cannot be forwarded due to improper system security settings or syslog itself. • The container running time, for example, shorter than 30s, is extremely short. • A single node generates logs at a high speed, exceeding the allowed transmit bandwidth or log collection speed. Ensure that the log generation speed of a single node is lower than 5 MB/s.
	Log loss	When a single log line exceeds 1024 bytes, this line will be discarded.
	Log repetition	When the ICAgent is restarted, identical data may be collected around the restart time.
Alarm	Alarm	You can query the alarms generated in the last 31 days.
	Event	You can query the events generated in the last 31 days.
-	Application discovery rule	You can create a maximum of 100 application discovery rules.

Service Usage Restrictions

If the AMS-Access service is powered off or restarted unexpectedly when you use AOM, a metric data breakpoint occurs on some resources such as hosts, components, and containers in a collection period. This breakpoint is visible on the monitoring page and has no impacts. To avoid breakpoints in a metric graph, set the value of **Interpolation Mode** to **0** or **average** on the **Metric Monitoring** page. In this way, the system automatically replaces breakpoints with **0** or average values, as shown in **Figure 8-1**.

Figure 8-1 Changing the interpolation mode

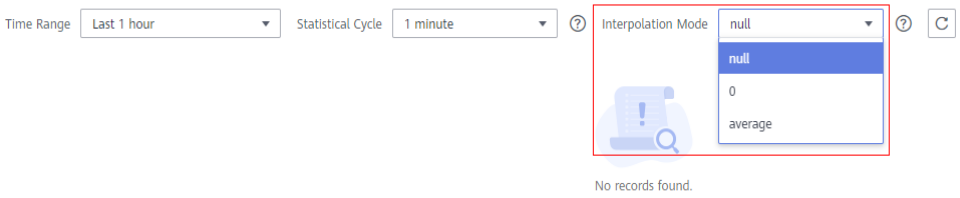

9 Privacy and Sensitive Information Protection Statement

All O&M data will be displayed on the AOM console. Therefore, do not upload your privacy or sensitive data to AOM. If necessary, encrypt such data.

Collector Deployment

When you manually install the ICAgent on an Elastic Cloud Server (ECS), your AK/SK will be used as an input parameter in the installation command. To prevent privacy leakage, disable historical record collection before installing the ICAgent. After the ICAgent is installed, it will encrypt and store your AK/SK.

Container Monitoring

For Cloud Container Engine (CCE) container monitoring, the AOM collector (ICAgent) must run as a privileged container. Evaluate the security risks of the privileged container and identify your container service scenarios. For example, for a node that provides services through logical multi-tenant container sharing, use open-source tools such as Prometheus to monitor the services and do not use ICAgent.

10 Relationships Between AOM and Other Services

AOM can work with Simple Message Notification (SMN), Distributed Message Service (DMS), and Cloud Trace Service (CTS). For example, when you subscribe to SMN, AOM can inform related personnel of threshold rule status changes by email or Short Message Service (SMS) message. When AOM interconnects with middleware services such as Virtual Private Cloud (VPC) and Elastic Load Balance (ELB), you can monitor them in AOM. When AOM interconnects with Cloud Container Engine (CCE) or Cloud Container Instance (CCI), you can monitor their basic resources and applications, and view related logs and alarms.

SMN

SMN can push notifications by SMS message, email, or app based on your requirements. You can integrate application functions through SMN to reduce system complexity.

AOM uses the message transmission mechanism of SMN. When it is inconvenient for you to query threshold rule status changes on site, AOM sends such changes to you by email or SMS messages. In this way, you can obtain resource status and other information in real time and take necessary measures to avoid service loss.

OBS

Object Storage Service (OBS) is a secure, reliable, and cost-effective cloud storage service. With OBS, you can easily create, modify, and delete buckets, as well as upload, download, and delete objects.

AOM allows you to dump logs to OBS buckets for long-term storage.

CTS

CTS records operations on cloud resources in your account. Based on the records, you can perform security analysis, monitor resource changes, conduct compliance audits, and locate faults. To store operation records for a longer time, you can subscribe to OBS and synchronize operation records to OBS in real time.

With CTS, you can record operations associated with AOM for future query, audit, and tracing.

IAM

Identity and Access Management (IAM) provides identity authentication, permission management, and access control.

IAM can implement authentication and fine-grained authorization for AOM.

Cloud Eye

Cloud Eye provides a multi-dimensional monitoring platform for resources such as Elastic Cloud Server (ECS) and bandwidth. With Cloud Eye, you can view the resource usage and service running status in the cloud, and respond to exceptions in a timely manner to ensure that services run smoothly.

By calling Cloud Eye APIs, AOM can obtain and display monitoring data of ECS, VPC, Relational Database Service (RDS), and Distributed Cache Service (DCS), so that you can monitor these services on the AOM console.

APM

APM monitors and manages the performance of cloud applications in real time. It provides performance analysis of distributed applications, helping O&M personnel quickly locate and resolve faults and performance bottlenecks.

AOM integrates APM functions to better monitor and manage applications.

VPC

VPC is a logically isolated virtual network. It is created for ECSs, and supports custom configuration and management, improving resource security and simplifying network deployment.

After subscribing to VPC, you can monitor VPC running status and metrics on the AOM console without installing other plug-ins.

ELB

ELB distributes access traffic to multiple backend ECSs based on forwarding policies. By distributing traffic, ELB expands the capabilities of application systems to provide services externally. By preventing single points of failures, ELB improves the availability of application systems.

After subscribing to ELB, you can monitor ELB running status and metrics on the AOM console without installing other plug-ins.

RDS

Relational Database Service (RDS) is a cloud-based web service that is reliable, scalable, and easy to manage.

After subscribing to RDS, you can monitor RDS running status and metrics on the AOM console without installing other plug-ins.

DCS

DCS is an online, distributed, in-memory cache service compatible with Redis, Memcached, and In-Memory Data Grid (IMDG). It is reliable, scalable, ready to

use out-of-the-box, and easy to manage, meeting your requirements for high read/write performance and fast data access.

After subscribing to DCS, you can monitor DCS running status and metrics on the AOM console without installing other plug-ins.

CCE

CCE is a high-performance and scalable container service through which enterprises can build reliable containerized applications. It integrates network and storage capabilities, and is compatible with Kubernetes and Docker container ecosystems. CCE enables you to create and manage diverse containerized workloads easily. It also provides efficient O&M capabilities, such as container fault self-healing, monitoring log collection, and auto scaling.

You can monitor basic resources, applications, logs, and alarms about CCE on the AOM console.

ServiceStage

ServiceStage is a one-stop PaaS platform service for enterprises. It hosts applications of enterprises on the cloud to simplify application lifecycle management, covering deployment, monitoring, O&M, and governance. In addition, ServiceStage provides a microservice framework compatible with mainstream open-source ecosystems and decoupled from specific development frameworks and platforms, helping enterprises quickly build distributed applications based on microservice architectures.

You can monitor basic resources, applications, logs, and alarms about ServiceStage on the AOM console.

FunctionGraph

FunctionGraph hosts and computes functions in a serverless context. It automatically scales up/down resources during peaks and spikes without requiring the reservation of dedicated servers or capacities. Resources are billed on a pay-per-use basis.

You can monitor basic resources, applications, logs, and alarms about FunctionGraph on the AOM console.

IEF

Intelligent EdgeFabric (IEF) provides you a complete edge computing solution, in which cloud applications are extended to the edge. By leveraging edge-cloud synergy, you can manage edge nodes and applications remotely and process data nearby, to meet your requirements for remote control, data processing, analysis, decision-making, and intelligence of edge computing resources. In addition, you can perform O&M in the cloud, including edge node monitoring, application monitoring, and log collection.

You can monitor resources (such as edge nodes, applications, and functions), logs, and alarms about IEF on the AOM console without installing other plug-ins.

ECS

ECS is a computing server consisting of the CPU, memory, image, and Elastic Volume Service (EVS) disk. It supports on-demand allocation and auto scaling. ECSs integrate VPC, virtual firewall, and multi-data-copy capabilities to create an efficient, reliable, and secure computing environment. This ensures stable and uninterrupted running of services. After creating an ECS server, you can use it like using your local computer or physical server.

When purchasing an ECS, ensure that its OS meets the requirements in **Table 8-1**. In addition, install an ICAgent on the ECS. Otherwise, the ECS cannot be monitored by AOM. You can monitor basic resources, applications, logs, and alarms about this ECS on the AOM console.

BMS

Bare Metal Server (BMS) is a dedicated physical server in the cloud. It provides high-performance computing and ensures data security for core databases, key application systems, and big data. With the advantage of scalable cloud resources, you can apply for BMS servers flexibly and they are billed on a pay-per-use basis.

When purchasing a BMS server, ensure that its OS meets the requirements in **Table 8-1**. In addition, install an ICAgent on the server. Otherwise, the server cannot be monitored by AOM. You can monitor basic resources, applications, logs, and alarms about this server on the AOM console.

11 Basic Concepts

Metrics

Metrics reflect resource performance data or status. A metric consists of a namespace, dimension, name, and unit.

Metric namespaces can be regarded as containers for storing metrics. Metrics in different namespaces are independent of each other so that metrics of different applications will not be aggregated to the same statistics information. Each metric has certain features, and a dimension may be considered as a category of such features. **Figure 11-1** describes the relationships among namespaces, dimensions, and cluster metrics.

Figure 11-1 Cluster metrics

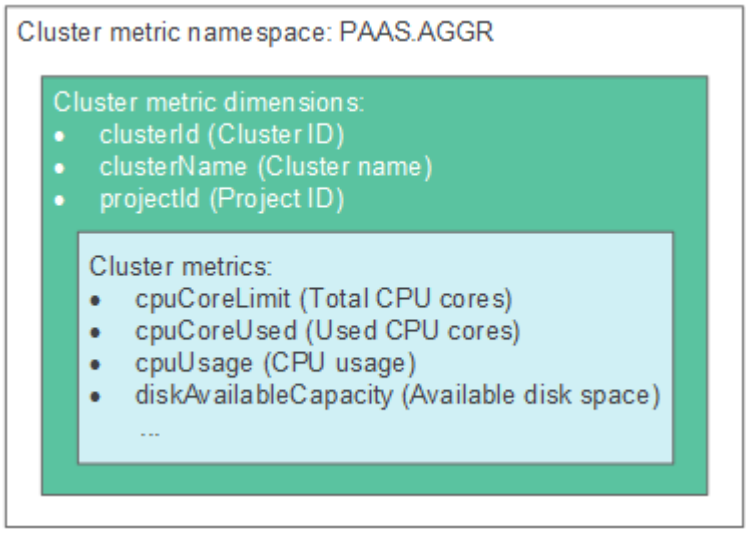

The metric storage duration and billing mode vary according to AOM editions. For details, see **AOM Pricing Details**.

Hosts

Each host of AOM corresponds to a VM or physical machine. A host can be your own VM or physical machine, or a VM (for example, an ECS) or physical machine

(for example, a BMS) that you purchased on Huawei Cloud. A host can only be connected to AOM for monitoring when its OS is supported by AOM and an ICAgent has been installed on the host. (For details about the OSs supported by AOM, see **OS Usage Restrictions**.)

ICAgent

ICAgent is the collector of AOM. It runs on hosts to collect metrics, logs, and application performance data in real time. Before using AOM, ensure that the ICAgent has been installed. Otherwise, AOM cannot be used.

Logs

AOM supports log collection, search, analysis, download, and dump. It also reports alarms based on keyword statistics and enables you to export reports, query SQL statements, and monitor data in real time.

The log storage duration, size, and billing mode vary according to AOM editions. For details, see **AOM Pricing Details**.

Log Buckets

Log buckets are logical groups of log files. Before creating statistical rules or querying bucket logs, ensure that a log bucket has been created.

Log Traffic

Log traffic refers to the volume of logs reported per second. A maximum of 10 MB/s is supported for each tenant in a region. If the log traffic exceeds 10 MB/s, logs may be lost.

Bucket Logs

Bucket logs support fine-grained query. You can view logs by bucket to obtain key service data, and quickly identify and locate problems.

Bucket logs allow you to query logs from multiple dimensions. You can also query and analyze original logs, and structured logs based on SQL syntax.

Alarms

Alarms are reported when AOM or an external service such as ServiceStage, Application Performance Management (APM), or Cloud Container Engine (CCE) is abnormal or may cause exceptions. Alarms will cause service exceptions and need to be handled.

There are two alarm clearance modes:

- Automatic clearance: After a fault is rectified, AOM automatically clears the corresponding alarm, for example, a threshold alarm.
- Manual clearance: After a fault is rectified, AOM does not automatically clear the corresponding alarm, for example, ICAgent installation failure alarm. In such a case, manually clear the alarm.

Events

Events generally carry some important information. They are reported when AOM or an external service, such as ServiceStage, APM, or CCE encounters some changes. Such changes do not necessarily cause service exceptions. Events do not need to be handled.

Threshold Rules

Threshold rules: You can set threshold conditions for resource metrics. AOM reports a threshold alarm when the value of a metric reaches the preset threshold, or reports an insufficient data event when no metric data is reported. In addition, a custom trigger policy is executed. When the threshold rule status (**Exceeded**, **OK**, or **Insufficient**) changes, a notification is sent by email or SMS message. In this way, you can detect and handle exceptions at the earliest time.

Notification Rules

AOM provides the notification function. When an alarm is reported due to an exception in AOM or an external service, alarm information will be sent to the specified personnel by email or SMS message. Therefore, such personnel can rectify faults in time to avoid service loss.

Statistical Rules

AOM can periodically count keywords or SQL statements and generate metric data, enabling you to monitor system performance and service information in real time. You can also set threshold rules for metrics. AOM reports a threshold alarm when the value of a metric reaches the preset threshold. In this way, you can detect and handle exceptions at the earliest time.

Topologies

Topologies show the call and dependency relationships between services. A topology consists of circles, lines with arrows, and resources. Each circle represents a service, and each segment in the circle represents an instance. The fraction in each circle indicates the number of active instance/total number of instances. The values below a fraction respectively indicate the number of calls, latency, and number of errors. Each line with an arrow represents a call relationship. Thicker lines indicate more calls. The values above a line respectively indicate the throughput and total latency. Throughput indicates the number of calls within the selected period. Application Performance Index (Apdex) is used in topologies to quantify user satisfaction with application performance. Different colors indicate different Apdex ranges, helping you quickly detect and locate faults.

Figure 11-2 Topology

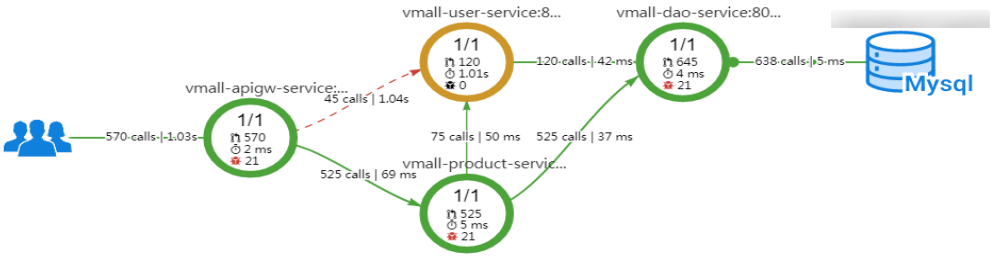

Transactions

In real life, a transaction is a one-time task. A user completes a task by using an application. For example, a product query in an e-commerce application is a transaction, and a payment is also a transaction. A transaction is usually an HTTP request (complete process: request > web server > database > web server > request).

Tracing

By tracing and recording service calls, AOM visually restores the execution traces and statuses of service requests in distributed systems, so that you can quickly locate performance bottlenecks and faults.

12 Permissions

If you need to assign different permissions to employees in your enterprise to access your AOM resources, Identity and Access Management (IAM) is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you secure access to your AOM resources.

With IAM, you can use your account to create IAM users for your employees, and assign permissions to the users to control their access to specific types of resources. For example, some software developers in your enterprise need to use AOM resources but are not allowed to delete them or perform any high-risk operations such as deleting application discovery rules. To achieve this result, you can create IAM users for the software developers and grant them only the permissions required for using AOM resources.

If your cloud account does not need individual IAM users for permissions management, you may skip over this chapter.

IAM can be used free of charge. You pay only for the resources in your account. For more information about IAM, see **IAM Service Overview**.

AOM Permissions

By default, new IAM users do not have any permissions assigned. You need to add a user to one or more groups, and assign permissions policies or roles to these groups. The user then inherits permissions from the groups it is a member of. This process is called authorization. After authorization, the user can perform specified operations on AOM.

AOM is a project-level service deployed and accessed in specific physical regions. To assign AOM permissions to a user group, specify the scope as region-specific projects and select projects for the permissions to take effect. If **All projects** is selected, the permissions will take effect for the user group in all region-specific projects. When accessing AOM, the users need to switch to a region where they have been authorized to use this service.

You can grant users permissions by using roles and policies.

- Roles: A type of coarse-grained authorization mechanism that defines permissions related to user responsibilities. This mechanism provides only a limited number of service-level roles for authorization. When using roles to

grant permissions, you also need to assign dependency roles. However, roles are not an ideal choice for fine-grained authorization and secure access control.

- Policies: A type of fine-grained authorization mechanism 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. For example, you can grant Elastic Cloud Server (ECS) users only the permissions for managing a certain type of ECSs. Most policies define permissions based on APIs. For the API actions supported by AOM, see **Permissions Policies and Supported Actions**.

Table 12-1 lists all the system permissions supported by AOM.

Table 12-1 System permissions supported by AOM

Policy Name	Description	Type	Depended System Permissions
AOM FullAccess	Administrator permissions for AOM. Users granted these permissions can operate and use AOM.	System-defined policy	CCE Administrator, OBS Administrator, and LTS FullAccess
AOM ReadOnlyAccess	Read-only permissions for AOM. Users granted these permissions can only view AOM data.	System-defined policy	

Table 12-2 lists the common operations supported by each system-defined policy of AOM. Please choose proper system-defined policies according to this table.

Table 12-2 Common operations supported by each system-defined policy of AOM

Operation	AOM FullAccess	AOM ReadOnlyAccess
Creating a threshold rule	√	x
Modifying a threshold rule	√	x
Deleting a threshold rule	√	x
Creating a threshold template	√	x
Modifying a threshold template	√	x
Deleting a threshold template	√	x
Creating a dashboard	√	x

Operation	AOM FullAccess	AOM ReadOnlyAccess
Modifying a dashboard	√	x
Deleting a dashboard	√	x
Creating an alarm action rule	√	x
Modifying an alarm action rule	√	x
Deleting an alarm action rule	√	x
Creating a message template	√	x
Modifying a message template	√	x
Deleting a message template	√	x
Creating a grouping rule	√	x
Modifying a grouping rule	√	x
Deleting a grouping rule	√	x
Creating a suppression rule	√	x
Modifying a suppression rule	√	x
Deleting a suppression rule	√	x
Creating a silence rule	√	x
Modifying a silence rule	√	x
Deleting a silence rule	√	x
Creating an application discovery rule	√	x
Modifying an application discovery rule	√	x
Deleting an application discovery rule	√	x
Exporting a monitoring report	√	√

Operation	AOM FullAccess	AOM ReadOnlyAccess
Configuring a VM log collection path	√	x
Viewing bucket logs	√	√
Adding a log dump	√	x
Modifying a log dump	√	x
Deleting a log dump	√	x
Starting periodical dump	√	x
Stopping periodical dump	√	x
Creating a statistical rule	√	x
Modifying a statistical rule	√	x
Deleting a statistical rule	√	x
Configuring a delimiter	√	x
Installing the ICAgent	√	√
Upgrading the ICAgent	√	x
Uninstalling the ICAgent	√	x

To use a custom fine-grained policy, log in to IAM as the administrator and select fine-grained permissions of AOM as required. For details, see **Table 12-3**

Table 12-3 Fine-grained permissions of AOM

Permission Name	Description	Dependency	Scenario
aom:alarmRule:create	Creating a threshold rule	N/A	Creating a threshold rule
aom:alarmRule:set	Modifying a threshold rule		Modifying a threshold rule
aom:alarmRule:get	Querying threshold rules		Querying all threshold rules or a single threshold rule by rule ID

Permission Name	Description	Dependency	Scenario
aom:alarmRule:delete	Deleting threshold rules		Deleting threshold rules in batches or a single threshold rule by rule ID
aom:discoveryRule:list	Querying application discovery rules		Querying existing application discovery rules
aom:discoveryRule:delete	Deleting an application discovery rule		Deleting an application discovery rule
aom:discoveryRule:set	Adding an application discovery rule		Adding an application discovery rule
aom:metric:list	Querying time series objects		Querying time series objects
aom:metric:list	Querying time series data		Querying time series data
aom:metric:get	Querying metrics		Querying metrics
aom:metric:get	Querying monitoring data		Querying monitoring data

Reference Links

- **IAM Service Overview**
- **Creating a User and Granting Permissions**
- **Permissions and Supported Actions**

13 Billing

Billing

AOM supports two editions: the basic edition and the pay-per-use edition. The pay-per-use edition includes postpaid and prepaid packages. Prepaid packages can be further divided into professional and enterprise packages. For more information, see **AOM Pricing Details**.

NOTE

- AOM interconnects with other cloud services to provide functions such as notification, log dump, and performance management. These functions may incur extra fees, which are settled according to standard pricing of corresponding cloud services.
  - Threshold rule and alarm notification: Based on Simple Message Notification (SMN), AOM sends the changes of threshold rule status and alarms to you by emails or Short Message Service (SMS) message. In this way, you can obtain information such as resource running status in real time and take necessary measures to avoid service loss. **SMN Pricing Details**
  - Log dump: Based on Object Storage Service (OBS), AOM dumps log files to OBS buckets for long-term storage. **OBS Pricing Details**
  - Log and threshold alarm subscription: Based on Distributed Message Service (DMS) for Kafka, AOM sends log or threshold alarm data to specified DMS Kafka queues, so that you can retrieve the data from these queues. **DMS for Kafka Pricing Details**
  - Application Performance Management (APM): Based on APM, AOM can provide more advanced O&M capabilities. **APM Pricing Details**

Switching Editions

You can switch AOM between the basic edition and the pay-per-use edition.

CAUTION

You can purchase packages only when you use the pay-per-use edition.

- Switching from the basic edition to the pay-per-use edition

  If the basic edition does not meet your requirements, switch to the pay-per-use edition. Specifically, click **Switch Editions** on the **Overview** page of AOM.

If you do not purchase any package after switching AOM to the pay-per-use edition, you will be billed on a pay-per-use basis for all the resources used. If you purchase a package and exceed the resource quota included in the package during its validity period, you will be billed on a pay-per-use basis for the excess resources used.

- Switching from the pay-per-use edition to the basic edition

  If you do not need the pay-per-use edition, click **Switch Editions** on the **Overview** page of AOM. Note that the basic edition can be rolled back only once every 24 hours. From 00:00 on the next day after AOM is switched to the basic edition, AOM deletes the resources that exceed the quota specified in the basic edition. Deleted resources cannot be recovered. Therefore, exercise caution when performing this operation.

  If you use the pay-per-use edition and your account is in arrears and frozen, your account will enter a retention period of 15 days. You can top up your account or manually switch AOM to the basic edition. If you manually switch AOM to the basic edition, AOM will delete the resources that exceed the quota specified in the basic edition. Deleted resources cannot be recovered. Therefore, exercise caution when performing this operation. If you do not top up your account, AOM will automatically switch to the basic edition after the retention period expires. From 00:00 on the next day after AOM is switched to the basic edition, AOM deletes the resources that exceed the quota specified in the basic edition. Deleted resources cannot be recovered. Therefore, top up your account in time.

Package Renewal

When your AOM package expires or is about to expire:

To continue using the package, renew it before it expires. If you do not renew the package before it expires, you will be billed on a pay-per-use basis on the next day after the package expires. If your account balance is insufficient to cover AOM resource fees, your account will be frozen and enter a retention period of 15 days. You can top up your account or manually switch AOM to the basic edition. If you manually switch AOM to the basic edition, AOM will delete the resources that exceed the quota specified in the basic edition. Deleted resources cannot be recovered. Therefore, exercise caution when performing this operation. If you do not top up your account, AOM will automatically switch to the basic edition after the retention period expires. From 00:00 on the next day after AOM is switched to the basic edition, AOM deletes the resources that exceed the quota specified in the basic edition. Deleted resources cannot be recovered. Therefore, top up your account in time.

14 Change History

Table 14-1 Change history

Released On	Description
2022-02-25	• Supported alarm rules. By setting alarms rules, you can define event conditions for services or threshold conditions for resource metrics. If the resource data of a service meets the event condition, an event alarm will be generated. If a metric value meets the threshold condition, a threshold alarm will be generated. If no metric data is reported, an insufficient data event will be generated. • Supported alarm action rules. You can associate an SMN topic with a message template for sending notifications. • Supported alarm noise reduction rules. Alarm noise reduction consists of four parts: grouping, deduplication, suppression, and silence. Alarms are processed based on the alarm noise reduction rules before notifications are sent.
2021-08-20	Connected AOM logs to LTS. By adding access rules, you can map logs of Cloud Container Engine (CCE) or custom clusters from AOM to LTS. Then you can view and analyze logs on LTS. Mapping does not generate extra fees, but duplicate mapping will.
2020-07-30	• Released AOM 2.0. • Switched to console UI 4.0.
2021-07-12	Modified metric names in **Disk Partition Metrics**.
2019-11-30	Supported cloud service monitoring. You can learn historical performance curves and running status of cloud service instances.
2019-01-31	Removed the function of configuring container log collection paths.
2018-12-05	Supported the function of configuring container log collection paths. **Learn more**
2018-11-26	Supported query of bucket logs. **Learn more**

Released On	Description
2018-10-31	Supported configuration of delimiters. You can separate log contents into multiple words by using delimiters, and then search for logs based on these words. **Learn more**
2018-10-24	Supported creation of statistical rules. AOM can periodically count keywords in log files and generate metrics data. **Learn more**
2018-09-26	Supported the log dump function. You can dump log files in log buckets to Object Storage Service (OBS) buckets. **Learn more**
2018-09-13	• Multi-dimensional cloud application O&M: Provides a full-link, multi-layer, and one-stop O&M platform for resources, applications, and application experience. • Device-side analysis: Supports performance metric and crash analysis of browser and mobile applications, achieving full control over applications. • Transaction insights: Supports automatic discovery of transaction performance problems, intelligent filtering, and root cause analysis (RCA). • Middleware monitoring: Supports monitoring of the statuses and metrics of middleware such as Relational Database Service (RDS) and Distributed Cache Service (DCS) on the AOM console without installing other plug-ins. • Supported creation of notification rules, enabling alarm information to be sent to specified personnel by Short Message Service (SMS) message or email. • Supported disk and file system monitoring.
2018-07-20	Provided both basic and professional editions, meeting your different demands for metric and log storage.
2018-06-12	This issue is the first official release.

