

Application Operations Management

User Guide

Date 2024-07-15

Contents

1 Service Overview.....	1
1.1 What Is AOM?.....	1
1.2 Product Architecture.....	3
1.3 Functions.....	3
1.4 Application Scenarios.....	5
1.5 Metric Overview.....	6
1.5.1 Introduction.....	6
1.5.2 Network Metrics and Dimensions.....	7
1.5.3 Disk Metrics and Dimensions.....	8
1.5.4 Disk Partition Metrics.....	9
1.5.5 File System Metrics and Dimensions.....	10
1.5.6 Host Metrics and Dimensions.....	11
1.5.7 Cluster Metrics and Dimensions.....	14
1.5.8 Container Metrics and Dimensions.....	16
1.5.9 VM Metrics and Dimensions.....	20
1.5.10 Instance Metrics and Dimensions.....	22
1.5.11 Service Metrics and Dimensions.....	22
1.6 Restrictions.....	22
1.7 Privacy and Sensitive Information Protection Statement.....	26
1.8 Relationships Between AOM and Other Services.....	26
1.9 Basic Concepts.....	29
1.10 Permissions.....	30
2 Getting Started.....	35
2.1 Process of Using AOM.....	35
2.2 Installing an ICAgent.....	36
3 Permissions Management.....	38
3.1 Creating a User and Granting Permissions.....	38
3.2 Creating a Custom Policy.....	39
4 Connecting Resources to AOM.....	41
4.1 Installing an ICAgent.....	41
4.2 Configuring Application Discovery Rules.....	43
4.3 Configuring VM Log Collection Paths.....	47

5 Monitoring Overview.....	50
5.1 O&M.....	50
5.2 Dashboard.....	53
6 Alarm Management.....	58
6.1 Alarm Rules.....	58
6.1.1 Overview.....	58
6.1.2 Alarm Tags and Annotations.....	58
6.1.3 Creating a Threshold Rule.....	59
6.1.4 Creating a Static Threshold Template.....	64
6.1.5 Creating an Event Alarm Rule.....	66
6.2 Checking Alarms.....	68
6.3 Checking Events.....	69
6.4 Alarm Action Rules.....	69
6.4.1 Overview.....	69
6.4.2 Creating an Alarm Action Rule.....	69
6.4.3 Creating a Message Template.....	71
6.5 Alarm Noise Reduction.....	74
6.5.1 Overview.....	74
6.5.2 Creating a Grouping Rule.....	76
6.5.3 Creating a Suppression Rule.....	79
6.5.4 Creating a Silence Rule.....	81
7 Resource Monitoring.....	83
7.1 Application Monitoring.....	83
7.2 Component Monitoring.....	84
7.3 Host Monitoring.....	86
7.4 Container Monitoring.....	88
7.5 Metric Monitoring.....	88
8 Log Management.....	90
8.1 Searching for Logs.....	90
8.2 Viewing Log Files.....	91
9 Configuration Management.....	93
9.1 ICAgent Management.....	93
9.1.1 Installing an ICAgent.....	93
9.1.2 Upgrading the ICAgent.....	97
9.1.3 Uninstalling the ICAgent.....	98
9.2 Log Configuration.....	100
9.2.1 Setting the Log Quota.....	100
9.2.2 Configuring Delimiters.....	100
9.2.3 Setting Log Collection.....	104
9.3 Quota Configuration.....	104
9.4 Metric Configuration.....	104

10 Auditing	105
10.1 Operations Logged by CTS	105
10.2 Querying Real-Time Traces	108
11 Upgrading to AOM 2.0	111
12 FAQs	113
12.1 User FAQs	113
12.2 Consultation FAQs	115
12.2.1 What Are the Usage Restrictions of AOM?	115
12.2.2 What Are the Differences Between AOM and APM?	119
12.2.3 How Do I Distinguish Alarms from Events?	119
12.2.4 What Is the Relationship Between the Time Range and Statistical Cycle?	120
12.2.5 Does AOM Display Logs in Real Time?	120
12.2.6 How Can I Do If I Cannot Receive Any Email Notification After Configuring a Threshold Rule?	121
12.2.7 Why Are Connection Channels Required?	121
12.3 Usage FAQs	121
12.3.1 What Can I Do If I Do Not Have the Permission to Access SMN?	121
12.3.2 What Can I Do If Resources Are Not Running Properly?	121
12.3.3 How Do I Set the Full-Screen Online Duration?	123
12.3.4 How Do I Obtain an AK/SK?	124
12.3.5 How Can I Check Whether a Service Is Available?	124
12.3.6 Why Is the Status of an Alarm Rule Displayed as "Insufficient"?	124
12.3.7 Why the Status of a Workload that Runs Normally Is Displayed as "Abnormal" on the AOM Page?	125
12.3.8 How Do I Create the apm_admin_trust Agency?	125
12.3.9 What Can I Do If an ICAgent Is Offline?	126
12.3.10 Why Is "no crontab for root" Displayed During ICAgent Installation?	127
A Change History	128


1 Service Overview

1.1 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



Multi-Dimensional O&M

Provides one-stop multi-dimensional O&M platform for mobile apps, networks, services, middleware, and cloud resources.

Figure 1-4 AOM advantage 2



Health Check

Monitors service health in real time and detects exceptions or performance bottlenecks within minutes.



Ease of Use

Connects to applications without having to modify codes and collects data in a non-intrusive way.

- **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.

1.2 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 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 metric data of time series, 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).

1.3 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 line graphs, 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 period. 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.

1.4 Application Scenarios

Problem Inspection and Demarcation

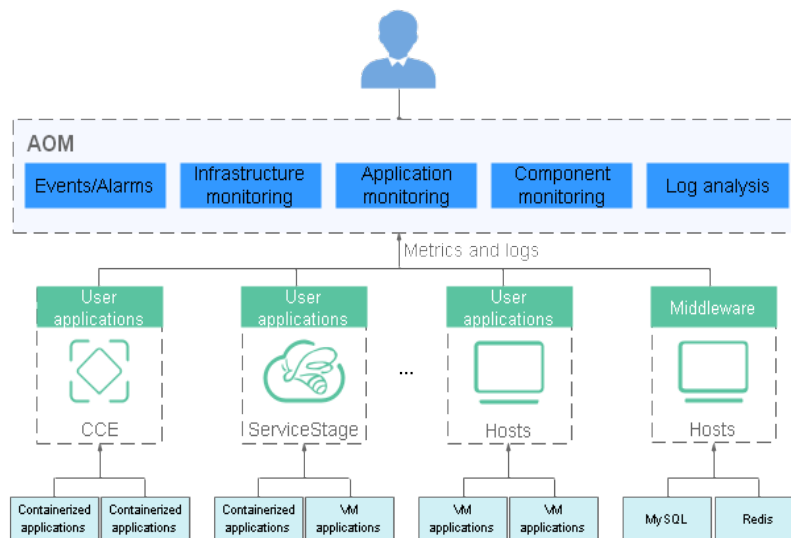
During routine O&M, it is hard to locate faults and obtain logs. Therefore, a monitoring platform is required to monitor resources, logs, and application performance.

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 1-5 Problem inspection and demarcation



Multi-Dimensional O&M

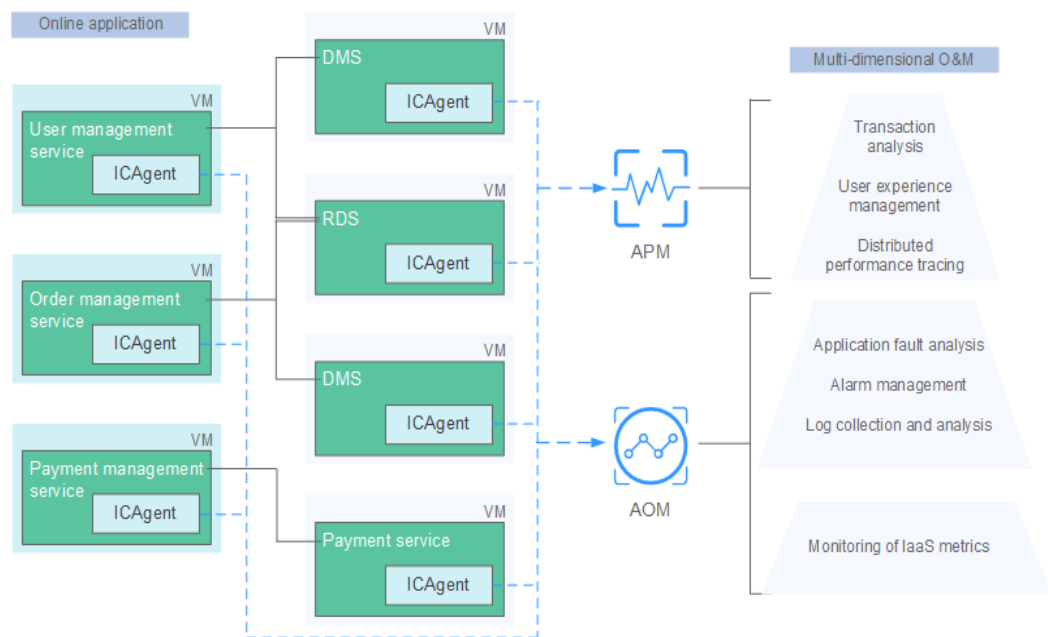
You need to monitor comprehensive system running status and make fast response to various problems.

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 1-6 Multi-dimensional O&M



1.5 Metric Overview

1.5.1 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" in the *Application Operations Management (AOM) API Reference*.
 - Method 2: When creating containerized applications on CCE, interconnect with Prometheus to report custom metrics. For details, see "Custom Monitoring" in *Cloud Container Engine (CCE) User Guide*.

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 1-1](#).

Table 1-1 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

- 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.

1.5.2 Network Metrics and Dimensions

Table 1-2 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

Metric	Description	Value Range	Unit
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 1-3 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

1.5.3 Disk Metrics and Dimensions

Table 1-4 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 1-5 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

1.5.4 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 1-6 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

1.5.5 File System Metrics and Dimensions

Table 1-7 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
Disk read/write status (aom_node_disk_rw_status)	Read or write status of a disk	0 or 1 <ul style="list-style-type: none"> • 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 1-8 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

1.5.6 Host Metrics and Dimensions

Table 1-9 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
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

Metric	Description	Value Range	Unit
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
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	<ul style="list-style-type: none"> • 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 <ul style="list-style-type: none"> • 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 <ul style="list-style-type: none"> • 0: Synchronous • 1: Not synchronized 	N/A

Metric	Description	Value Range	Unit
Processes (aom_node_process_number)	Number of processes on a measured object	≥ 0	N/A
GPU temperature (aom_node_gpu_temperature_centrigrade)	GPU temperature of a measured object	-	°C
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 1-10 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

1.5.7 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 1-11 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

Metric	Description	Value Range	Unit
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
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	%

Metric	Description	Value Range	Unit
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
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 1-12 Dimensions of cluster metrics

Dimension	Description
clusterId	Cluster ID
clusterName	Cluster name
projectId	Project ID

1.5.8 Container Metrics and Dimensions

Table 1-13 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

Metric	Description	Value Range	Unit
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
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

Metric	Description	Value Range	Unit
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	%
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

Metric	Description	Value Range	Unit
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
Status (aom_process_status)	Docker container status	0 or 1 <ul style="list-style-type: none"> • 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 1-14 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

Dimension	Description
kind	Application type
nameSpace	Cluster namespace
podID	Instance ID
podName	Instance name
serviceID	Inventory ID
gpuID	GPU ID
npuName	NPU name
npuID	NPU ID

1.5.9 VM Metrics and Dimensions

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

Table 1-15 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

Metric	Description	Value Range	Unit
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
Status (aom_process_status)	Process status	0 or 1 <ul style="list-style-type: none"> • 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 1-16 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

1.5.10 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 [1.5.8 Container Metrics and Dimensions](#) and [1.5.9 VM Metrics and Dimensions](#).

1.5.11 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 [1.5.10 Instance Metrics and Dimensions](#).

1.6 Restrictions

OS Usage Restrictions

AOM supports multiple operating systems (OSs). When creating a host, ensure that its OS meets the requirements in [Table 1-17](#). Otherwise, the host cannot be monitored by AOM.

Table 1-17 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 1-18](#).

Table 1-18 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	<ul style="list-style-type: none"> • 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	Metric data can be stored in the database for up to 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.
	Metric query API	A maximum of 20 metrics can be queried at a time.
	Statistical period	The maximum statistical period is 1 hour.

Category	Object	Usage Restrictions
	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	<ul style="list-style-type: none"> 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.
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.
	Log file	Only text log files can be collected. Other types of log files, such as binary files, cannot be collected.
<p>The ICAgent can collect a maximum of 20 log files from a volume mounting directory.</p> <p>The ICAgent can collect a maximum of 1000 standard container output log files. These files must be in JSON format.</p>		

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: <ul style="list-style-type: none"> • 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 15 days.
	Event	You can query the events generated in the last 15 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.

1.7 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.

1.8 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.

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.

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.

RDS

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

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.

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 1-17](#). 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 1-17](#). 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.

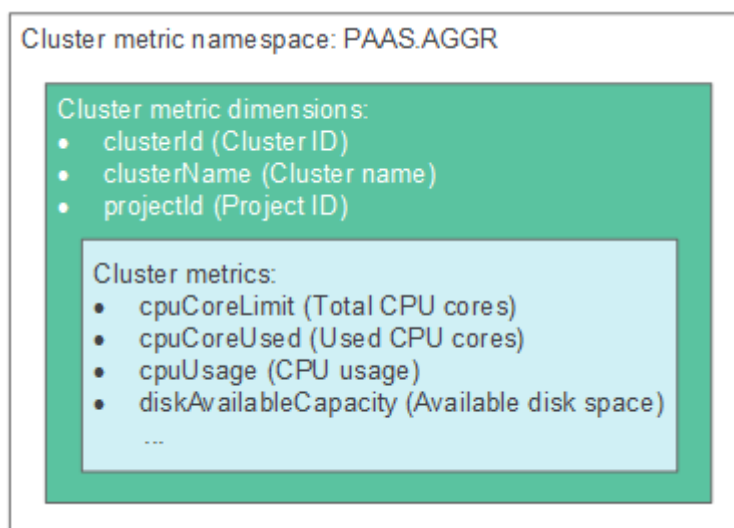
1.9 Basic Concepts

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 1-7](#) describes the relationships among namespaces, dimensions, and cluster metrics.

Figure 1-7 Cluster metrics



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) that you created. 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.

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.

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.

1.10 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.

Table 1-19 lists all the system permissions supported by AOM.

Table 1-19 System permissions supported by AOM

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

Policy Name	Description	Type	Depended System Permissions
AOM Viewer	Read-only permissions for AOM. Users granted these permissions can only view AOM data.	System-defined policy	

Table 1-20 lists the common operations supported by each system-defined policy of AOM. Please choose proper system-defined policies according to this table.

Table 1-20 Common operations supported by each system-defined policy of AOM

Operation	AOM Admin	AOM Viewer
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
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

Operation	AOM Admin	AOM Viewer
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	√	√
Configuring a VM log collection path	√	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 1-21](#)

Table 1-21 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
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

2 Getting Started

2.1 Process of Using AOM

Application Operations Management (AOM) is a one-stop cloud operations management platform that helps you with problem management, monitoring, security, and performance. Streamline cloud operational processes and effectively manage cloud hardware, software, services, and networks.

Figure 2-1 Process of using AOM



1. (Optional) Create IAM users and set permissions.

Create IAM users for employees based on the organizational structure of your enterprise, and grant different access permissions to them.

2. (Mandatory) Create a cloud host.

A host corresponds to a VM, for example, Elastic Cloud Server (ECS), or a physical machine, for example, Bare Metal Server (BMS). A host can be directly created on the ECS or BMS console, or indirectly created on the Cloud Container Engine (CCE) console.

3. (Mandatory) **Install an ICAgent.**

An ICAgent is a collector of AOM. It collects metrics, logs, and application performance data in real time. For hosts created on the ECS or BMS console, install the ICAgent manually. For hosts created on the CCE console, the ICAgent is automatically installed.

4. (Optional) Configure application discovery rules.

Connect applications on the host to AOM for monitoring. For the applications that meet built-in application discovery rules, they will be automatically discovered after the ICAgent is installed. For those that do not meet built-in rules, custom your own rules.

5. (Optional) Configure a log collection path.

To use AOM to monitor host logs, configure a log collection path first.

6. (Optional) Perform O&M.

Use AOM functions such as , and alarm management to perform routine O&M.

2.2 Installing an ICAgent

This section describes how to install an ICAgent on an ECS and monitor its resources on the AOM console.

Prerequisites

- You have created an ECS.
- You have **obtained an AK/SK**.
- The time of the local browser must be consistent with that of the ECS.

Installing an ICAgent

Step 1 Log in to the AOM console. In the navigation pane, choose **Configuration Management > Agent Management**.

Step 2 Select **Other: custom hosts**, and click **Install ICAgent**.

Step 3 Click **Copy Command** to copy the installation command.

Step 4 Log in to the ECS remotely.

Specifically, log in to the ECS console, click **Remote Login** in the **Operation** column of the target ECS, and then log in to the ECS as the **root** user.

Step 5 Run the ICAgent installation command.

On the ECS page, click **Copy & Paste**. On the page that is displayed, press **Ctrl+V** to paste the ICAgent installation command obtained in **Step 3**, and click **Send** to

send the command to the CLI. In the CLI, press **Enter** to run the ICAgent installation command.

If the message "ICAgent install success" is displayed, the ICAgent is successfully installed in the **/opt/oss/servicemgr/** directory. Wait for a while and go back to the **Agent Management** page to check whether the ICAgent status of the ECS is **Running**.

- If the ICAgent status is **Running**, the ICAgent is successfully installed.
- If the ICAgent status is **Offline**, view details. The AK/SK or ECS agency may be incorrect. In this case, obtain the correct AK/SK or reconfigure the ECS agency, and then reinstall the ICAgent.

Step 6 On the AOM console, monitor the ECS.

After the ICAgent is installed, wait for about 1–2 minutes. Then, in the navigation pane, choose **Overview > O&M** to monitor the ECS.

----End

3 Permissions Management

3.1 Creating a User and Granting Permissions

This chapter describes how to use Identity and Access Management (IAM) for fine-grained permissions control for your AOM resources. With IAM, you can:

- Create IAM users for employees based on the organizational structure of your enterprise. Each IAM user has their own security credentials, providing access to AOM resources.
- Grant only the permissions required for users to perform a task.
- Entrust a cloud account or service to perform professional and efficient O&M on your AOM resources.

If your account does not need individual IAM users, then you may skip over this section.

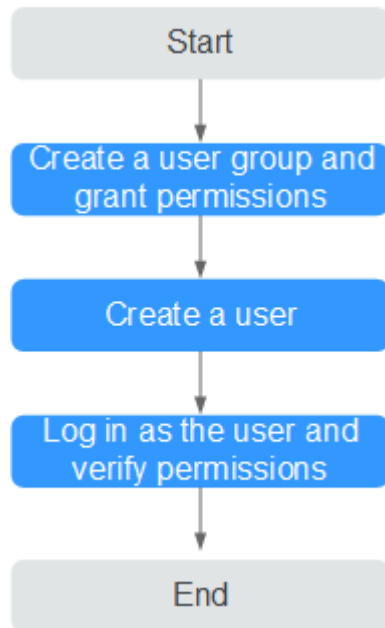
This section describes the procedure for granting permissions (see [Figure 3-1](#)).

Prerequisites

Learn about the permissions supported by AOM and choose policies or roles according to your requirements. See section "Permissions Management" in *AOM Service Overview*. For the permissions of other services, see section "Permission Description" in the help center.

Process

Figure 3-1 Process for granting AOM permissions



1. Create a user group and assign permissions.
Create a user group on the IAM console, and assign the **AOM ReadOnlyAccess** policy to the group.
2. Create an IAM user.
Create a user on the IAM console and add the user to the group created in **1**.
3. Log in and verify permissions.
Log in to the AOM console as the created user, and verify that it only has read permissions for AOM.

3.2 Creating a Custom Policy

Custom policies can be created as a supplement to the system policies of AOM.

You can create custom policies in either of the following two ways:

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.
- JSON: Edit JSON policies from scratch or based on an existing policy.

For example custom policies, see the following description.

Example Custom Policies

- Example 1: Allowing a user to create threshold rules

```
{  
  "Version": "1.1",  
  "Statement": [  
    {  
      "Effect": "Allow",
```

```

    "Action": [
      "aom:alarmRule:create"
    ]
  }
]
}

```

- Example 2: Forbidding a user to delete application discovery rules

A deny policy must be used in conjunction with other policies to take effect. If the permissions assigned to a user contain both Allow and Deny actions, the Deny actions take precedence over the Allow actions.

To grant a user the **AOM FullAccess** system policy but forbid the user to delete application discovery rules, create a custom policy that denies the deletion of application discovery rules, and grant both the **AOM FullAccess** and deny policies to the user. Because the Deny action takes precedence, the user can perform all operations except deleting application discovery rules. The following is an example deny policy:

```

{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": [
        "aom:discoveryRule:delete"
      ]
    }
  ]
}

```

- Example 3: Defining permissions for multiple services in a policy

A custom policy can contain actions of multiple services that are all of the project-level type. The following is an example policy containing actions of multiple services:

```

{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "aom*:list",
        "aom*:get",
        "apm*:list",
        "apm*:get"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "cce:cluster:get",
        "cce:cluster:list",
        "cce:node:get",
        "cce:node:list"
      ]
    }
  ]
}

```

4 Connecting Resources to AOM

4.1 Installing an ICAgent

ICAgents collect metrics, logs, and application performance data in real time. For hosts purchased from the ECS or BMS console, you need to manually install the ICAgent. For hosts purchased from the CCE console, the ICAgent is automatically installed.

Prerequisites

- Before installing an ICAgent, ensure that the time and time zone of the local browser are consistent with those of the server. If multiple servers are deployed, ensure that the local browser and multiple servers use the same time zone and time. Otherwise, metric data of applications and servers displayed on the UI may be incorrect.
- An ICAgent process needs to be installed and run by the **root** user.

Installation Methods

There are two methods to install an ICAgent. Note that the two methods are not applicable to container nodes created through ServiceStage or CCE. For container nodes, you do not need to manually install an ICAgent. Instead, you only need to perform certain operations when creating clusters or deploying applications.

For details, see [Table 4-1](#).

Table 4-1 Installation methods

Method	Scenario
Initial installation	This method is used when the following condition is met: An ICAgent has never been installed on your server.

Method	Scenario
Inherited installation	This method is used when the following conditions are met: You have multiple servers where an ICAgent is to be installed. One server is bound to an EIP, but others are not. An ICAgent has been installed on the server bound to an EIP by using the initial installation method. You can use the inherited method to install an ICAgent on the remaining servers. See Inherited Installation .

Initial Installation

After you apply for a server and install an ICAgent for the first time, perform the following operations:

Step 1 Obtain an Access Key ID/Secret Access Key (AK/SK).

- If you have obtained the AK/SK, skip this step.
- If you have not obtained an AK/SK, [obtain them first](#).

Step 2 In the navigation pane, choose **Configuration Management > Agent Management**.

Step 3 Select **Other: custom hosts**, and click **Install ICAgent**.

Step 4 Click **Copy Command** to copy the installation command.

Step 5 Use a remote login tool to log in to the target server as the **root** user, and run the following command to disable historical record collection:

```
set +o history
```

Step 6 Run the copied installation command and enter the obtained AK and SK as prompted.

Step 7 After the ICAgent is installed, run the following command to enable historical record collection:

```
set -o history
```

NOTE

- If the message **ICAgent install success** is displayed, the ICAgent has been installed in the **/opt/oss/servicemgr/** directory. After the ICAgent has been installed, choose **Configuration Management > Agent Management** in the navigation pane to view the ICAgent status.
- If the ICAgent fails to be installed, uninstall the ICAgent according to [Uninstalling the ICAgent by Logging In to the Server](#) and then install it again. If the problem persists, contact technical support.

----End

Follow-up Operations

For more information about how to install, upgrade, and uninstall the ICAgent, see [9.1 ICAgent Management](#).

4.2 Configuring Application Discovery Rules

AOM can discover applications and collect their metrics based on configured rules. There are two modes to configure application discovery: auto mode and manual mode. This section mainly describes the manual mode.

- **Auto mode**

After you install the ICAgent on a host according to [9.1.1 Installing an ICAgent](#), the ICAgent automatically discovers applications on the host based on [Built-in Discovery Rules](#) and displays them on the **Application Monitoring** page.

- **Manual mode**

If you customize an application discovery rule and apply it to the host where the ICAgent is installed (for details, see [9.1.1 Installing an ICAgent](#)), the ICAgent discovers applications on the host based on the custom rule and displays them on the **Application Monitoring** page.

Filtering Rules

The ICAgent will periodically implement detection on the target host to find out all its processes. The effect is similar to that of running the **ps -e -o pid,comm,lstart,cmd | grep -v defunct** command on the target host. Then, the ICAgent checks whether processes match the filtering rules in [Table 4-2](#). If a process meets a filtering rule, the process is filtered out and is not discovered by AOM. If a process does not meet any filtering rules, the process is not filtered out and is discovered by AOM.

Information similar to the following is displayed:

PID	COMMAND	STARTED	CMD
1	systemd	Tue Oct 2 21:12:06 2018	/usr/lib/systemd/systemd --switched-root --system --deserialize 20
2	kthreadd	Tue Oct 2 21:12:06 2018	[kthreadd]
3	ksoftirqd/0	Tue Oct 2 21:12:06 2018	(ksoftirqd/0)
1140	tuned	Tue Oct 2 21:12:27 2018	/usr/bin/python -Es /usr/sbin/tuned -l -P
1144	sshd	Tue Oct 2 21:12:27 2018	/usr/sbin/sshd -D
1148	agetty	Tue Oct 2 21:12:27 2018	/sbin/agetty --keep-baud 115200 38400 9600 hvc0 vt220
1154	docker-containe	Tue Oct 2 21:12:29 2018	docker-containerd -l unix:///var/run/docker/libcontainerd/docker-containerd.sock --shim docker-containerd-shim --start-timeout 2m --state-dir /var/run/docker/libcontainerd/containerd --runtime docker-runc --metrics-interval=0

Table 4-2 Filtering rules

Filtering Rule	Example
If the COMMAND value of a process is docker-containe, vi, vim, pause, sshd, ps, sleep, grep, tailf, tail, or systemd-udevd , and the process is not running in the container, the process is filtered out and is not discovered by AOM.	In the preceding information, the process whose PID is 1154 is not discovered by AOM because its COMMAND value is docker-containe .
If the CMD value of a process starts with [and ends with] , the process is filtered out and is not discovered by AOM.	In the preceding information, the process whose PID is 2 is not discovered by AOM because its CMD value is [kthreadd] .
If the CMD value of a process starts with (and ends with) , the process is filtered out and is not discovered by AOM.	In the preceding information, the process whose PID is 3 is not discovered by AOM because its CMD value is (ksoftirqd/0) .
If the CMD value of a process starts with /sbin/ , the process is filtered out and is not discovered by AOM.	In the preceding information, the process whose PID is 1148 is not discovered by AOM because its CMD value starts with /sbin/ .

Built-in Discovery Rules

AOM provides two built-in discovery rules: **Sys_Rule** and **Default_Rule**. These rules are executed on all hosts, including hosts added later. The priority of **Sys_Rule** is higher than that of **Default_Rule**. That is, **Sys_Rule** is executed on the host first. If **Sys_Rule** is met, **Default_Rule** is not executed. Otherwise, **Default_Rule** is executed. Rule details are as follows:

Sys_Rule (cannot be disabled)

When **Sys_Rule** is used, the component name and application name must be used together. The names are determined according to the following priorities:

- Priorities for determining the application name:
 - a. Use the value of the **Damp_application** field in the process startup command.
 - b. If the value in **a** is empty, use the value of the **Dapm_application** field in the **JAVA_TOOL_OPTIONS** variable.
 - c. If the value in **b** is empty, use the value of the **PAAS_MONITORING_GROUP** variable.
 - d. If the value in **c** is empty, use the value of the **DAOM.APPN** field in the process startup command.
- Priorities for determining the component name:
 - a. Use the value of the **DAOM.PROCN** field in the process startup command. If the value is empty, use the value of the **Dapm_tier** field.

- b. If the value in **a** is empty, use the value of the **Dapm_tier** field in the **JAVA_TOOL_OPTIONS** variable.
- c. If the value in **b** is empty, use the value of the **PAAS_APP_NAME** variable.

In the following example, the component name is **atps-demo** and the application name is **atpd-test**.

```
PAAS_MONITORING_GROUP=atpd-test  
PAAS_APP_NAME=atps-demo  
JAVA_TOOL_OPTIONS=-javaagent:/opt/oss/servicemgr/ICAgent/pinpoint/pinpoint-bootstrap.jar -  
Dapm_application=atpd-test -Dapm_tier=atps-demo
```

Default_Rule (can be disabled)

- If the **COMMAND** value of a process is **java**, obtain the name of the JAR package in the command, the main class name in the command, and the first keyword that does not start with a hyphen (-) in the command based on the priorities in descending order as the component name, and use the default value **unknownapplicationname** as the application name.
- If the **COMMAND** value of a process is **python**, obtain the name of the first .py/.pyc script in the command as the component name, and use the default value **unknownapplicationname** as the application name.
- If the **COMMAND** value of a process is **node**, obtain the name of the first .js script in the command as the component name, and use the default value **unknownapplicationname** as the application name.

Custom Discovery Rules

Step 1 In the navigation pane, choose **Configuration Management > Application Discovery**.

Step 2 Click **Add Custom Application Discovery Rule** and configure an application discovery rule.

Step 3 Select a host for pre-detection.

1. Customize a rule name, for example, **rule-test**.
2. Select a typical host, for example, **host-test**, to check whether the application discovery rule is valid. The hosts that execute the rule will be configured in **Step 6**. Then, click **Next**.

Step 4 Set an application discovery rule.

1. Click **Add Check Items**. AOM can discover processes that meet the conditions of check items.

For example, AOM can detect the processes whose command parameters contain **ovs-vsitchd unix:** and environment variables contain **SUDO_USER=paas**.

NOTE

- To precisely detect processes, you are advised to add check items about unique features of the processes.
 - You must add at least one check item and can add up to five check items. If there are multiple check items, AOM only discovers the processes that meet the conditions of all check items.
2. After adding check items, click **Detect** to search for the processes that meet the conditions.

If no process is detected within 20s, modify the discovery rule and detect processes again. Only when at least one process is detected can you proceed to the next step.

Step 5 Set an application name and component name.

Set an application name.

1. Set an application name.

In the **Application Name Settings** area, click **Add Naming Rule** to set an application name for the detected process.

 **NOTE**

- If you do not set an application name, the default name **unknownapplicationname** is used.
- When you add multiple naming rules, all the naming rules are combined as the application name of the process. Metrics of the same application are aggregated.


2. Set a component name.

In the **Component Name Settings** area, specify an application type and click **Add Naming Rule** to set a component name for the discovered process. For example, add the text **app-test** as a component name.

 **NOTE**

- Application types are specified to identify application categories. They are used only for better rule classification and console display. You can enter any field. For example, you can enter **Java** or **Python** to categorize applications by technology stack or enter **collector** or **database** to categorize applications by function.
- If you do not set a component name, the default name **unknownapplicationname** is used.
- When you add multiple naming rules, all the naming rules are combined as the component name of the process. Metrics of the same component are aggregated.

3. Preview the component name.

If the name does not meet your requirements, click  in the **Preview Component Name** table to rename the component.

Step 6 Set a priority and detection range.

1. Set a priority: When there are multiple rules, set priorities. Enter 1 to 9999. A smaller value indicates a higher priority. For example, **1** indicates the highest priority and **9999** indicates the lowest priority.
2. Set a detection range: Select a host to be detected. That is, select the host to which the configured rule is applied. If no host is selected, this rule will be executed on all hosts, including hosts added later.

Step 7 Click **Add** to complete the configuration. AOM collects metrics of the process.

Step 8 After about two minutes, choose **Monitoring > Component Monitoring** in the navigation pane, select the target host from the cluster drop-down list, and find out the monitored component.

----End

More Operations

After creating an application discovery rule, perform the operations listed in [Table 4-3](#) if needed.

Table 4-3 Related operations

Operation	Description
Viewing rule details	In the Name column, click the name of an application discovery rule.
Enabling or disabling a rule	<ul style="list-style-type: none"> Click Enable in the Operation column. Click Disable in the Operation column. After a rule is disabled, AOM does not collect corresponding process metrics.
Deleting a rule	<ul style="list-style-type: none"> To delete a discovery rule, click Delete in the Operation column. To delete one or more application discovery rules, select them and click Delete above the rule list. <p>NOTE Built-in application discovery rules cannot be deleted.</p>
Modifying a rule	<p>Click Modify in the Operation column.</p> <p>NOTE Built-in application discovery rules cannot be modified.</p>

4.3 Configuring VM Log Collection Paths

AOM can collect and display VM logs. VM refers to an Elastic Cloud Server (ECS) or a Bare Metal Server (BMS) running Linux. To use this function, first configure a log collection path according to the following procedure.

Prerequisites

- You have installed an ICAgent on a VM according to [Installing an ICAgent](#). Wait for about 5 minutes after the installation is complete. Then you can view the VM in the VM list on the **Path Configuration** page.

Precautions

- Ensure that your VMs are ECSs or BMSs running Linux.
- If you specify a directory, all **.log**, **.trace**, and **.out** text log files in this directory are collected by default. If you specify a log file, only this file is collected. The specified file must be a text file. Other types of log files, such as binary log files, cannot be collected.
- Ensure that an absolute path of a log directory or file is configured and the path exists. For example, **/opt/yilu/work/xig** or **/opt/yilu/work/xig/debug_cpu.log**.

- The ICAgent does not collect log files from subdirectories. For example, the ICAgent does not collect log files from the `/opt/yilu/work/xig/debug` subdirectory of `/opt/yilu/work/xig`.
- A maximum of 20 log collection paths can be configured for a VM.
- If the difference between the last modification time of a log file and the current time exceeds 12 hours, the log file will not be collected.
- For ECSs in the same resource set, logs will be collected based on the latest log collection configuration. AOM and LTS log collection configurations cannot take effect at the same time. For example, if you set log collection paths in AOM for ECSs, the previous LTS collection configurations of all ECSs under the resource set become invalid.

Configuring Log Collection Paths for a Single VM Through the Console


Step 1 Log in to the AOM console. In the navigation pane, choose **Log > Path Configuration**. The **Host Log** tab page is displayed.

Step 2 In the VM list, click **Configure** in the **Operation** column to configure one or more log collection paths for a VM.


You can use the paths automatically identified by the ICAgent or manually configure paths.

- **Using the Paths Automatically Identified by the ICAgent**

The ICAgent automatically scans the log files of your VM, and displays all the `.log`, `.trace`, or `.out` log files with handles and their paths on the page.

You can click  in the **Operation** column to add a path automatically identified by the ICAgent to the log collection path list. To configure multiple paths, repeat this operation.

- **Manually Configuring Log Collection Paths**

If the paths automatically identified by the ICAgent cannot meet your requirements, enter a log directory or file (such as `/opt/yilu/work/xig/debug_cpu.log` or `/opt/yilu/work/xig/*.log`) in the **Log Collection Path** text box, and then click  to add the path to the log collection path list. To configure multiple paths, repeat this operation.

Step 3 Click **OK**.

----End

Configuring Log Collection Paths for Multiple VMs in Batches Through the Console

You can configure log collection paths for multiple VMs in batches. When your service is deployed on multiple VMs, you can configure log collection paths in batches to reduce workload.

Step 1 Log in to the AOM console. In the navigation pane, choose **Log > Path Configuration**. The **Host Log** tab page is displayed.

Step 2 Configure one or more log collection paths for multiple VMs in batches.

Select one or more VMs in the list, click **Batch Configure**, and enter a log directory or file (for example, `/opt/yilu/work/xig/debug_cpu.log`) in the **Log Collection Path** text box.

 **NOTE**

If you configure log collection paths for your VM and then configure log collection paths in batches, new paths will be added to the existing path list.

Step 3 Click **OK**.

In the VM list, click  in the **Log Collection Path** column to view the configured log collection paths of the VM.

----End

Viewing VM Logs

After the log collection paths are configured, the ICAgent collects log files from them. This operation takes about 1 minute to complete. After collecting logs, you can perform the following operations:

- **Viewing VM Log Files**

In the navigation pane, choose **Log > Log Files**. Click the **Host** tab to view the collected log files. For details, see [8.2 Viewing Log Files](#).

- **Viewing and Analyzing VM logs**

In the navigation pane, choose **Log > Log Search**. Click the **Host** tab to view and analyze the collected logs by time range, keyword, and context. For details, see [8.1 Searching for Logs](#).

5 Monitoring Overview

5.1 O&M

The **O&M** page supports full-link, multi-layer, and one-stop O&M for resources, applications, and user experience. Specifically, this page displays the following types of cards: infrastructure monitoring, application monitoring, alarm statistics, host monitoring (CPU and memory), component monitoring (CPU and memory), container instance monitoring (CPU and memory), host monitoring (disk), host monitoring (network), cluster monitoring (CPU and memory), and cluster monitoring (disk) cards.

Infrastructure Monitoring Card

This card mainly displays infrastructure metrics. You can select one or all clusters to view information. When you select all clusters, the following information is displayed:

- Host running status, CPU usage, and physical memory usage.
- Trend graph of network traffic in the last 30 minutes. The values of each point in the graph respectively indicate the total receive rate (BPS) and send rate (BPS) of all clusters in one minute. The values above the graph respectively indicate the total receive rate (BPS) and send rate (BPS) of all clusters at the latest time point.
- Trend graph of CPU and memory usage in the last 30 minutes. The values of each point in the graph respectively indicate the average CPU and memory usage of all clusters in one minute. The values above the graph respectively indicate the average CPU and memory usage of all clusters at the latest time point.

Application Monitoring Card

This card mainly displays application metrics:

1. Running status of applications, components, containers, and instances.
2. When you select an application, the following information is displayed:
 - Trend graph of network traffic in the last 30 minutes. The values of each point in the graph respectively indicate the receive rate (BPS) and send

rate (BPS) of the selected application in one minute. The values above the graph respectively indicate the receive rate (BPS) and send rate (BPS) of the selected application at the latest time point.

- Trend graph of CPU and memory usage in the last 30 minutes. The values of each point in the graph respectively indicate the CPU and memory usage of the selected application in one minute. The values above the graph respectively indicate the CPU and memory usage of the selected application at the latest time point.

Alarm Statistics Card

This card mainly displays alarms, log usage, threshold rules, and trends of alarms and hosts.

Component Monitoring (CPU and Memory) Card

This card mainly displays:

- The top 5 components with high CPU and memory usage in the last minute.
- Trend graph of the CPU and memory usage of the selected component in the last hour. The values of each point in the graph respectively indicate the CPU and memory usage of the component in one minute.
- CPU and memory usage of the selected component at the latest time point, which is displayed above the trend graph.
- You can select **Hide system components** in the lower left corner.

Cluster Monitoring (Disk) Card

This card mainly displays:

- The top 5 clusters with high disk usage in the last minute.
- Trend graph of the disk usage of the selected cluster in the last hour. The value of each point in the graph indicates the disk usage of the cluster in one minute.
- Disk usage of the selected cluster at the latest time point, which is displayed above the trend graph.

Container Instance Monitoring (CPU and Memory) Card

This card mainly displays:

- The top 5 container instances with high CPU and memory usage in the last minute.
- Trend graph of the CPU and memory usage of the selected container instance in the last hour. The values of each point in the graph respectively indicate the CPU and memory usage of the container instance in one minute.
- CPU and memory usage of the selected container instance at the latest time point, which is displayed above the trend graph.
- You can select **Hide system instances** in the lower left corner.

Host Monitoring (Disk) Card

This card mainly displays:

- The top 5 hosts with high disk read/write rate in the last minute.
- Trend graph of the disk read/write rate of the selected host in the last hour. The values of each point in the graph respectively indicate the disk read/write rate of the selected host in one minute.
- Disk read/write rate of the selected host at the latest time point, which is displayed above the trend graph.

Host Monitoring (Network) Card

This card mainly displays:

- The top 5 hosts with high send/receive rate in the last minute.
- Trend graph of the send/receive rate of the selected host in the last hour. The values of each point in the graph respectively indicate the send/receive rate of the selected host in one minute.
- Send/receive rate of the selected host at the latest time point, which is displayed above the trend graph.

Host Monitoring (CPU and Memory) Card

This card mainly displays:

- The top 5 hosts with high CPU and memory usage in the last minute.
- Trend graph of the CPU and memory usage of the selected host in the last hour. The values of each point in the graph respectively indicate the CPU and memory usage of the host in one minute.
- CPU and memory usage of the selected host at the latest time point, which is displayed above the trend graph.

Cluster Monitoring (CPU and Memory) Card




This card mainly displays:

- The top 5 clusters with high CPU and memory usage in the last minute.
- Trend graph of the CPU and memory usage of the selected cluster in the last hour. The values of each point in the graph respectively indicate the CPU and memory usage of the cluster in one minute.
- CPU and memory usage of the selected cluster at the latest time point, which is displayed above the trend graph.

More Operations

Perform the operations listed in [Table 5-1](#) if needed.

Table 5-1 Related operations

Operation	Description
Adding a card to favorites	To hide a card, click  in the upper right corner of the card and choose Add to Favorites . After a card is added to favorites, it is hidden from the O&M page. To view the card later, obtain it from favorites.
Adding a card to dashboard	Click  in the upper right corner of the card and choose Add to Dashboard .
Zooming in a metric graph	Click  in the upper right corner of the metric graph.
Drilling down blue texts	Click the blue texts, such as Host , Application , or Component to drill down to the details page.

5.2 Dashboard

With a dashboard, different graphs such as line graphs and digit graphs are displayed on the same screen, which lets you view comprehensive monitoring data.

For example, you can add key metrics of important resources to the dashboard for real-time monitoring. You can also compare the same metric for different resources on one GUI. In addition, you can add routine O&M metrics to the dashboard so that you can perform routine check without re-selecting metrics when you re-open AOM.

Before creating a dashboard, learn the types of graphs that can be added to the dashboard for accurate resource monitoring. The following graphs can be added to the dashboard:

Metric Data Graphs (Including Line and Digit Graphs)

- **Line graph:** displays the metric data trend by time. Use this type of graph to monitor the metric data trend of one or more resources in a period.
You can use a line graph to compare the same metric of different resources.
- **Digit graph:** displays the latest value of a metric in real time.

Health Status Graphs (Including Threshold, Host, and Component Status Graphs)

The status of thresholds, hosts, and components can be displayed. The status of one or more threshold rules, hosts, or components can be added to one graph for monitoring.

- **Threshold-crossing status graph:** monitors the status of threshold rules in real time.

 **NOTE**

Before adding a threshold status graph, ensure that you have [created a threshold rule](#). Otherwise, such a graph cannot be added.

- **Host status graph:** monitors the host status in real time.
- **Component status graph:** monitors the component status in real time.

Top N Resource Graphs

For top N resource graphs, the statistical unit is a cluster and statistical objects are resources such as hosts, components, or instances in the cluster. A top N resource graph shows the top N resources in a cluster in a visualized manner. Both the top 5 and top 15 resources can be displayed. By default, the top 5 resources are displayed. After the graph is zoomed in, the top 15 resources are displayed.

To quickly view the top N resources, add a top N graph to the dashboard. You only need to select resources and metrics, for example, host CPU usage. AOM then automatically singles out top N hosts for display. If the number of resources is less than N, actual resources are displayed.

 **NOTE**

- By default, the top 5 resources are displayed. To view the top 15 resources, click **Top 15 xxx**, double-click the graph, or click **View Larger** in the **Operation** column.
- To monitor the top 5 resources among all clusters, view them on the **O&M** page. Alternatively, add the corresponding graph on the **O&M** page to the dashboard.
- You can customize the title of the top N resource graph. By default, the title is **resource type(cluster name)**.

Precautions

- A maximum of 50 dashboards can be created in a region.
- A maximum of 20 graphs can be added to a dashboard.
- A maximum of 10 resources can be added to a line graph, and resources can be selected across clusters.
- Only one resource can be added to a digit graph.
- A maximum of 10 threshold rules can be added to a threshold-crossing 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.

Creating a Dashboard

Step 1 In the navigation pane, choose **Overview > Dashboard**.

Step 2 On the **Dashboard** page, click **Create Dashboard**. In the displayed **Create Dashboard** dialog box, enter a dashboard name and click **OK**.


Step 3 Add a metric graph to the dashboard. The dashboard supports the following graphs: line graphs, digit graphs, threshold-crossing status graphs, host status graphs, and component status graphs. Select a graph that meets your requirements.

The following shows how to add a line graph to a dashboard:

1. On the **Dashboard** page, click **Add Metric Graph**. In the displayed **Select Which to Add** dialog box, click **Create** below **Metric Data**.
2. Select the type of the graph: In the displayed **Add Metric Graph** dialog box, select **Line graph** and then click **Next**.
3. Select the metrics and set **Statistical Mode** and **Statistical Cycle**, and click **OK**.

Step 4 Click **Save** in the upper right corner of the **Dashboard** page.

 **NOTE**

Enable **Auto Refresh** () in the upper right corner of the **Dashboard** page so that all graphs in the dashboard can be refreshed automatically.

- On (default)
Data in the dashboard will be automatically refreshed each minute.
- Off
Data in the dashboard will not be automatically refreshed.



----End

More Operations

After creating a dashboard, perform the operations listed in [Table 5-2](#) if needed.

Table 5-2 Related operations

Object	Operation	Description
Dashboard	Save as	Click More in the upper right corner, and choose Save As , Rename , or Delete from the drop-down list.
	Rename	
	Delete	
	Export a monitoring report	Click Export Monitoring Report to export a line graph in the dashboard as a CSV file to a local PC.

Object	Operation	Description
	Set the full-screen online duration	<ol style="list-style-type: none"> 1. Select the target dashboard and click  in the upper right corner of the Dashboard page. 2. In the dialogue box that is displayed, set the full-screen online duration. <p>NOTE</p> <ul style="list-style-type: none"> • Custom: The default online duration is 1 hour. You can enter 1-24 (unit: hour) in the text box. For example, if you enter 2 in the text box, the login page is automatically displayed 2 hours later. • Always online: The full-screen online duration is not restricted. That is, you can always implement full-screen monitoring and the login page will never be displayed. • Rotation Period: Set Rotation Period and Dashboard if rotation display is enabled. Range: 10s (default) to 120s. <ol style="list-style-type: none"> 3. Click OK to enter the full-screen mode.
	Set an interpolation mode	<p>Click Interpolation Mode to set a mode for aggregating metric data. By default, AOM uses null to represent breakpoints in a metric graph. However, a metric graph with breakpoints is not suitable for reporting or presentation. To solve the problem, set Interpolation Mode to 0 or null to interpolate values. In this way, you can replace the missing metric data and avoid breakpoints.</p> <p>You can set Interpolation Mode to null, or 0.</p> <ul style="list-style-type: none"> • null: Breakpoints are represented by null by default. • 0: Breakpoints are indicated by 0.
Graph	Add	Click Add Metric Graph to add a line graph, digit graph, threshold-crossing status graph, host status graph, or component status graph to the dashboard.
	Edit	<p>Choose Edit, Copy, Delete, and View Larger (only a line graph can be enlarged) from the Operation column. The Time Select option is available only in a line graph. This option allows you to set a temporary time range and statistical cycle so that you can view the resource data within a specified time range.</p> <p>NOTE</p> <p>In the dashboard, when resources such as hosts and components are deleted, graphs created for these resources are not automatically deleted. To improve system performance, manually delete unnecessary graphs.</p>
	Copy	
	Delete	
	Zoom in	
	Time select	
	Refresh	
Resize	Move the cursor to the lower right corner of a graph. When the cursor changes to  , hold down your left mouse button to resize the graph.	

Object	Operation	Description
	Reposition	Put the cursor at the blank area in the upper or lower part of a graph, and drag and drop it to the desired position.

6 Alarm Management

6.1 Alarm Rules

6.1.1 Overview

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 a threshold condition, a threshold alarm will be reported. If there is no metric data, an insufficient data event will be reported.

Alarm rules are classified into threshold rules and event alarm rules. Generally, threshold rules monitor the usage of resources such as hosts and components in real time. If there are too many resource usage alarms and notifications are sent too often, use an event alarm rule to identify a type of resource usage problems for simplified notification.

The total number of threshold rules and event alarm rules is 1000. If the number of alarm rules has reached the upper limit, delete unnecessary rules and create new ones.

6.1.2 Alarm Tags and Annotations

When creating alarm rules, you can set alarm tags and annotations. Tags are attributes that can be used to identify alarms. They are applied to the alarm noise reduction scenario. Annotations are attributes that cannot be used to identify alarms. They are applied to scenarios such as alarm notification and message templates.

Alarm Tags

- Alarm tags can apply to grouping rules, suppression rules, and silence rules. The alarm management system manages alarms and notifications based on the tags.
- Each tag is in "key:value" format and can be customized. The key and value can contain only letters, digits, and underscores (_) and cannot start with an underscore (_). You can create up to 10 custom tags.

- If you set a tag when creating an alarm rule, the tag is automatically added as an alarm attribute when an alarm is triggered.
- In the message template, the `$event.metadata.key1` variable specifies the tag. For details, see [Table 6-11](#).

Alarm Annotations

- Annotations are attributes that cannot be used to identify alarms. They are applied to scenarios such as alarm notification and message templates.
- Each annotation is in "key:value" format and can be customized. The key and value can contain only letters, digits, and underscores (_) and cannot start with an underscore (_). You can create up to 10 custom annotations.
- In the message template, the `$event.annotations.key2` variable specifies the annotation. For details, see [Table 6-11](#).

6.1.3 Creating a Threshold Rule

You can set threshold conditions for resource metrics by setting threshold rules. If a metric value meets a threshold condition, a threshold alarm will be reported. If there is no metric data, an insufficient data event will be reported.

Creation Methods

There are two creation methods: [Directly Creating Threshold Rules](#) and [Using Templates to Create Threshold Rules](#). Only one rule is generated at a time. All resources are monitored using the same rule. To use the second method to create threshold rules, ensure that a static threshold template has been created according to [6.1.4 Creating a Static Threshold Template](#).

Precautions

- If you need AOM to send email or SMS notifications when the threshold rule status (**Exceeded**, **Normal**, **Insufficient**, or **Disabled**) changes, set an alarm action rule according to [6.4.2 Creating an Alarm Action Rule](#).
- If you use a threshold rule to monitor the same metric of multiple resources in batches, pay attention to the following:
 - If the metric status of a resource is **Exceeded**, the status of the threshold rule is also **Exceeded**.
 - If the metric status of one or more resources is **Insufficient** or **Normal**, the status of the threshold rule is **Normal**.

Directly Creating Threshold Rules


Step 1 Log in to the AOM console. In the navigation pane, choose **Alarm Center > Alarm Rules**. Then, click **Create Alarm Rule** in the upper right corner.

Step 2 Set a threshold rule.

1. Set basic information such as the rule name and description.
2. Set details about the rule.
 - a. Set **Rule Type** to **Threshold alarm**.

- b. Select monitored objects. Use either of the following methods:
 - Select resource objects: Click **Select Resource Object**, add objects by dimension or resource, and click **Confirm**.

 **NOTE**


- A threshold rule can monitor up to 100 pieces of metric data.
- If you enable **Apply to All** () when selecting objects to monitor, an alarm rule will be created for all metrics of the type you select under an application or service. For example, if you select **CCE/Host/Host/CPU Usage** and enable **Apply to All**, an alarm rule will be created for all hosts in CCE.
- Click **Edit resource objects** to modify the selected resource object.

- Command input: Both manual and auto inputs are supported.


- Manual input: used when you know the metric name and IP address, and you are familiar with the Prometheus format.

For example, to query the CPU usage of the host, run command `avg(label_replace(avg_over_time(aom_node_cpu_usage{hostID="81010a40-1682-41c1-9645-f0588ff9c0cf",nodeIP="192.168.1.210",clusterId='00000000-0000-0000-0000-00000000'}[59999ms]), "__name__", "aom_node_cpu_usage", "", "")) by(__name__,hostID,nodeIP)`.

 **NOTE**

For details about Prometheus commands, move the cursor to  next to the search box and click [Learn more](#).

- Auto input: used when you do not know the metric information or are unfamiliar with the Prometheus format. The command can only be automatically filled when you switch from the **Metric Monitoring** page.

Specifically, choose **Monitoring > Metric Monitoring** in the navigation pane. Then, click **Add Metric** and select **Dimension** or **Resource** for **Add By**. Select up to 12 metrics to monitor. Next, click  in the **Operation** column. The system automatically switches to the threshold rule creation page and fills the Prometheus command for your metric.

- c. Set an alarm condition. Click **Custom** and set information such as **Statistical Period**, **Consecutive Periods**, and **Threshold Criterion**. [Table 6-1](#) describes the parameters.

Table 6-1 Alarm condition parameters

Category	Parameter	Description
Trigger Condition	Statistical Period	Interval at which metric data is collected. By default, only one period is measured. A maximum of five periods can be measured.

Category	Parameter	Description
	Consecutive Periods	When the metric value meets the threshold condition for a specified number of consecutive periods, a threshold-crossing alarm will be generated.
	Statistic	Method used to measure metrics. Options: Avg., Min., Max., Sum, and Samples.
	Threshold Condition	Trigger condition of a threshold alarm. A threshold condition consists of two parts: operators (\geq , \leq , $>$, and $<$) and threshold value. For example, after Threshold Criterion is set to > 85 , if the actual metric value exceeds 85, a threshold alarm is generated. Move the cursor to the graph area above the alarm condition. The ID, IP address, and unit of the current metric are displayed.
	Alarm Severity	Severity of a threshold alarm. Options: Critical, Major, Minor, and Warning.
Advanced Settings	Alarm Clearance	An alarm will be cleared if the monitored object does not meet the trigger condition within the monitoring period. By default, metrics in only one period are monitored. You can set up to five monitoring periods.
	Action Taken for Insufficient Data	Action to be taken when no metric data is generated or metric data is insufficient within the monitoring period. You can set this option based on your requirements. By default, metrics in only one period are monitored. You can set up to five monitoring periods. Options: Alarm, Insufficient data, Keep previous status, and Normal.


- d. Set alarm tags and annotations to group alarms. They can be associated with alarm noise reduction policies for sending alarm notifications. For details, see [Alarm Tags and Annotations](#).
Click **Add Tag** or **Add Annotation**.
- 3. Set an alarm notification policy. There are two alarm notification modes.
 - **Direct Alarm Reporting:** An alarm is directly sent when the alarm condition is met.
 - i. Specify whether to enable an alarm action rule. After an alarm action rule is enabled, the system sends notifications based on the associated SMN topic and message template. If the existing alarm

action rules cannot meet your requirements, click **Create Rule** to create one. For details, see [Creating an Alarm Action Rule](#).

- ii. After an alarm action rule is selected, specify whether to enable alarm clearance notification. After alarm clearance notification is enabled, if the alarm clearance condition set in [Advanced Settings > Alarm Clearance](#) is met, alarm clearance notifications are sent based on the selected action rule.

- **Alarm Noise Reduction:** Alarms are sent only after being processed based on noise reduction rules, preventing alarm storms.

Select a grouping rule from the drop-down list. If existing grouping rules cannot meet your requirements, click **Create Rule** to create one. For details, see [6.5.2 Creating a Grouping Rule](#).

Step 3 Click **Create Now**. A threshold rule is created. Click  to monitor the same metric of multiple resources.

In the expanded list, if the metric data of a host meets the preset alarm condition, a threshold alarm is generated on the alarm page. To view the alarm, go to the AOM console and choose **Alarm Center > Alarm List** in the navigation pane.

----End

Using Templates to Create Threshold Rules

Before creating threshold rules, ensure that a static threshold template has been created according to [6.1.4 Creating a Static Threshold Template](#).

Step 1 Log in to the AOM console. In the navigation pane, choose **Alarm Center > Alarm Rules**. Then, click **Create Alarm Rule** in the upper right corner.

Step 2 Set a threshold rule.


1. Set basic information such as the rule name and description.
2. Set details about the rule.
 - a. Set **Rule Type** to **Threshold alarm**.
 - b. Select monitored objects. When a template is used to create a threshold rule, you can select metrics only by dimension or resource. The command input mode is not supported.
 - c. Set an alarm condition. Click **Template**, select the created static threshold template from the drop-down list, and set parameters, such as **Alarm Clearance** and **Action Taken for Insufficient Data**.

Table 6-2 Alarm condition parameters

Category	Parameter	Description
Alarm Template	-	Select the static threshold template you have created. If the existing templates do not meet your requirements, click Create Alarm Template to create one. For details, see 6.1.4 Creating a Static Threshold Template .

Category	Parameter	Description
Trigger Condition	-	The system automatically imports the preset trigger condition in the template. Note that the condition cannot be modified.
Advanced Settings	Alarm Clearance	An alarm will be cleared if the monitored object does not meet the trigger condition within the monitoring period. By default, metrics in only one period are monitored. You can set up to five monitoring periods.
	Action Taken for Insufficient Data	Action to be taken when no metric data is generated or metric data is insufficient within the monitoring period. You can set this option based on your requirements. By default, metrics in only one period are monitored. You can set up to five monitoring periods. Options: Alarm , Insufficient data , Keep previous status , and Normal .

- d. Set alarm tags and annotations to group alarms. They can be associated with alarm noise reduction policies for sending alarm notifications.
Click **Add Tag** or **Add Annotation**.
3. Set an alarm notification policy. There are two alarm notification modes.
 - **Direct Alarm Reporting:** An alarm is directly sent when the alarm condition is met.
 - i. Specify whether to enable an alarm action rule. After an alarm action rule is enabled, the system sends notifications based on the associated SMN topic and message template. If the existing alarm action rules cannot meet your requirements, click **Create Rule** to create one. For details, see [Creating an Alarm Action Rule](#).
 - ii. After an alarm action rule is selected, specify whether to enable alarm clearance notification. After alarm clearance notification is enabled, if the alarm clearance condition set in [Advanced Settings > Alarm Clearance](#) is met, alarm clearance notifications are sent based on the selected action rule.
 - **Alarm Noise Reduction:** Alarms are sent only after being processed based on noise reduction rules, preventing alarm storms.
Select a grouping rule from the drop-down list. If existing grouping rules cannot meet your requirements, click **Create Rule** to create one. For details, see [6.5.2 Creating a Grouping Rule](#).

Step 3 Click **Create Now**. A threshold rule is created. Click  to monitor the same metric of multiple resources.


In the expanded list, if the metric data of a host meets the preset alarm condition, a threshold alarm is generated on the alarm page. To view the alarm, go to the AOM console and choose **Alarm Center > Alarm List** in the navigation pane.

----End

More Operations

After creating a threshold rule, perform the operations listed in [Table 6-3](#) if needed.

Table 6-3 Related operations

Operation	Description
Editing a threshold rule	Click Edit in the Operation column.
Deleting threshold rules	<ul style="list-style-type: none"> To delete a threshold rule, click Delete in the Operation column. To delete one or more threshold rules, select the check boxes before them and click Batch Operation and then Delete above the rule list.
Starting or stopping a threshold rule	<p>Click More > Start or Stop in the Operation column. Alternatively, select the check boxes before one or more threshold rules and choose Batch Operation > Start or Stop above the rule list.</p> <p>NOTE Single-resource threshold rules cannot be started or stopped.</p>
Searching for a threshold rule	You can search for a rule by rule name, description, or metric name. Simply enter a keyword in the search box in the upper right corner and click  .
Viewing an alarm	<p>When the metric value of a resource meets the threshold condition during the configured consecutive periods, the system reports a threshold alarm.</p> <p>In the navigation pane, choose Alarm Center > Alarm List to view the alarm.</p>
Viewing an event	<p>When no metric data of a resource is reported during the configured consecutive periods, the system reports an insufficient data event.</p> <p>In the navigation pane, choose Alarm Center > Event List to view the event.</p>

6.1.4 Creating a Static Threshold Template

Ensure that a static threshold template is available if you want to create threshold rules based on a template.

Precautions

You can create a maximum of 50 static threshold templates. If the maximum number has been reached, delete unnecessary templates and create new ones.

Procedure

Step 1 Log in to the AOM console. In the navigation pane, choose **Alarm Center > Alarm Rules**.

Step 2 Click the **Static Threshold Templates** tab, and then click **Create**.

Step 3 Customize a static threshold template.

Enter a template name, select a resource type, and set parameters such as **Name**, **Statistic**, and **Threshold Criterion**.

NOTE

- **Statistic:** method used to measure metric values.
- **Threshold Criterion:** trigger condition of a threshold alarm. It consists of two parts: determination condition (\geq , \leq , $>$, or $<$) and threshold value. For example, after **Threshold Criterion** is set to > 85 , if the actual metric value exceeds 85, a threshold alarm is generated.
- **Consecutive Periods:** If a metric value meets the threshold condition for a specified number of consecutive periods, a threshold alarm is generated.
- **Statistical Period:** interval at which metric data is collected.
- **Alarm Severity:** includes **Critical**, **Major**, **Minor**, and **Warning**.

Step 4 Click **Create**.


----End

More Operations

After creating a static threshold template, perform the operations listed in [Table 6-4](#) if needed.

Table 6-4 Related operations

Operation	Description
Using a static threshold template to create a multi-resource threshold rule	Click Create Rule in the Operation column. For details, see Using Templates to Create Threshold Rules .
Editing a static threshold template	Click Edit in the Operation column.

Operation	Description
Deleting a static threshold template	<ul style="list-style-type: none"> To delete a static threshold template, click Delete in the Operation column. To delete one or more static threshold templates, select them and click Delete above the template list.
Searching for a static threshold template	Enter a template name in the search box in the upper right corner and click  .

6.1.5 Creating an Event Alarm Rule

You can set event conditions for services by setting event alarm rules. When the resource data meets an event condition, an event alarm is generated.

Precautions

If you want to receive email or SMS notifications when the resource data meets the event condition, set an alarm action rule according to [6.4.2 Creating an Alarm Action Rule](#).

Procedure

- Step 1** Log in to the AOM console. In the navigation pane, choose **Alarm Center > Alarm Rules**. Then, click **Create Alarm Rule** in the upper right corner.
- Step 2** Set an event alarm rule.
1. Set basic information such as the rule name and description.
 2. Set details about the rule.
 - a. Set **Rule Type** to **Event alarm**.
 - b. Set the alarm source, trigger object, and trigger policy.

Table 6-5 Alarm rule parameters

Parameter	Description
Alarm Source	Name of the service for which an event alarm is reported. You can select a service from the service list.
Trigger Object	Select criteria to filter service events. You can select Notification Type , Event Name , Alarm Severity , Custom Attributes , Namespace , or Cluster Name as the filter criterion. One or more criteria can be selected.

Parameter	Description
Trigger Policy	<p>Policy for triggering event alarms.</p> <ul style="list-style-type: none"> ▪ Accumulated Triggering: An alarm action rule is triggered when the accumulated number of times you preset is reached in a monitoring period. ▪ Immediate Triggering: An alarm is generated immediately when the filter criterion is met.

3. Set an alarm notification policy. There are two alarm notification modes.

- **Direct Alarm Reporting:** An alarm is directly sent when the alarm condition is met.

You need to configure whether to enable an alarm action rule. After this function is enabled, the system sends alarm notifications based on the associated SMN topic and message template. If the existing alarm action rules cannot meet your requirements, click **Create Rule** to create one. For details, see [Creating an Alarm Action Rule](#).

- **Alarm Noise Reduction:** Alarms are sent only after being processed based on noise reduction rules, preventing alarm storms.

Select a grouping rule from the drop-down list. If existing grouping rules cannot meet your requirements, click **Create Rule** to create one. For details, see [6.5.2 Creating a Grouping Rule](#).

Step 3 Click **Create Now**.

This rule monitors critical alarm events of AOM. When a service event meets the preset notification policy, the system sends an alarm notification to specified personnel by email or SMS.

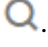
----End

Related Operations

After creating an event alarm rule, perform the operations listed in [Table 6-6](#) if needed.

Table 6-6 Related operations

Operation	Description
Editing an event alarm rule	Click Edit in the Operation column.
Deleting event alarm rules	<ul style="list-style-type: none"> • To delete an event alarm rule, click Delete in the Operation column. • To delete one or more event alarm rules, select the check boxes before them and click Delete above the rule list.


Operation	Description
Starting or stopping an event alarm rule	Click Start or Stop in the Operation column.
Searching for an event alarm rule	You can search for a rule by rule name, description, or metric name. Simply enter a keyword in the search box in the upper right corner and click  .

6.2 Checking Alarms

Procedure


Step 1 In the navigation pane, choose **Alarm Center > Alarm List**.

Step 2 Check alarms on the **Alarm List** page.

- Set a time range to check alarms. There are two methods to set a time range:
 Method 1: Use the predefined time label, such as **Last 1 hour**, **Last 6 hours**, or **Last 1 day**. Select one as required.
 Method 2: Specify the start time and end time to customize a time range. You can specify up to 31 days.
- Set filter criteria and click  to check the alarms generated in the period.

Step 3 Perform the operations listed in [Table 6-7](#) as required.

Table 6-7 Operations

Operation	Method	Description
Checking alarm statistics	Click Show Graph , and check alarm statistics that meet filter criteria within a specific time range on a bar graph.	-
Clearing alarms	In the alarm list, click  in the Operation column of the target alarm.	<ul style="list-style-type: none"> You can clear alarms after the problems that cause them are resolved. You can check the alarms that have been cleared on the History tab page.
Checking alarm details	Check alarm details in the Alarm Detail column.	-

----End


6.3 Checking Events

Events generally carry some important information, informing you of the changes of AOM or an external service. Such changes do not necessarily cause exceptions. Events do not need to be handled.

Procedure

Step 1 In the navigation pane, choose **Alarm Center > Event List**.

Step 2 Check events on the **Event List** page.

1. Set a time range to check events. There are two methods to set a time range:
 Method 1: Use the predefined time label, such as **Last 1 hour**, **Last 6 hours**, or **Last 1 day**. Select one as required.
 Method 2: Specify the start time and end time to customize a time range. You can specify up to 31 days.
2. Set filter criteria and click  to check the events generated in the period.

Step 3 Perform the operations listed in [Table 6-8](#) as required.

Table 6-8 Operations

Operation	Method	Description
Checking event statistics	Click Show Graph , and check event statistics that meet filter criteria within a specific time range on a bar graph.	-

----End

6.4 Alarm Action Rules

6.4.1 Overview

AOM allows you to customize alarm action rules. You can create an alarm action rule to associate an SMN topic with a message template. You can also customize notification content based on a message template. After an alarm action rule is created, choose **Alarm Center > Alarm Noise Reduction** in the navigation pane. Then, click the **Grouping Rules** tab and click **Create**. On the displayed page, specify an alarm action rule.

6.4.2 Creating an Alarm Action Rule

Prerequisites

- A topic has been created.

- A topic policy has been set.
- **APM** has been selected for **Services that can publish messages to this topic**. If **APM** is not selected, notifications cannot be sent.
- A subscriber, that is, an email or SMS message recipient has been added to the topic.

Precaution

You can create a maximum of 1000 alarm action rules. If this number has been reached, delete unnecessary rules.

Procedure

Step 1 Log in to the AOM console. In the navigation pane, choose **Alarm Center > Alarm Action Rules**. On the displayed page, click **Create** in the upper left corner.

Step 2 Set parameters such as **Rule Name** and **Action Type**.

Table 6-9 Parameters for configuring an alarm action rule

Parameter	Description
Rule Name	Name of an action rule. Enter 1 to 100 characters, and do not start or end with an underscore (_). Only digits, letters, and underscores are allowed.
Description	Description of the action rule.
Action Type	Type of action that is associated with the SMN topic and message template. Select your desired action from the drop-down list. Currently, only Notification is supported.
Topic	SMN topic. Select your desired topic from the drop-down list. If there is no topic you want to select, create one on the SMN console.
Message Template	Notification message template. Select your desired template from the drop-down list. If there is no message template you want to select, create one by referring to 6.4.3 Creating a Message Template .


Step 3 Click **OK**.

----End

More Operations

After an alarm action rule is created, you can perform operations described in [Table 6-10](#).

Table 6-10 Related operations

Operation	Description
Modifying an alarm action rule	Click Edit in the Operation column.
Deleting an alarm action rule	<ul style="list-style-type: none"> To delete a single rule, click Delete in the Operation column in the row that contains the rule, and then click Yes on the displayed page. To delete one or more rules, select them, click Delete above the rule list, and then click Yes on the displayed page. <p>NOTE Before deleting an alarm action rule, you need to delete the alarm rule bound to the action rule.</p>
Searching for an alarm action rule	Enter a rule name in the search box in the upper right corner and click  .

6.4.3 Creating a Message Template

You can create message templates to customize notifications. When a preset notification rule is triggered, notifications can be sent to specified personnel by emails, SMS, voice calls, HTTP, or HTTPS. If no message template is created, the default message template will be used.

Creating a Message Template

Step 1 Log in to the AOM console and choose **Alarm Center > Alarm Action Rules** in the navigation pane. On the displayed page, click the **Message Templates** tab.

Step 2 On the **Message Templates** page, click **Create**.

1. Enter a template name.
2. Enter a template description.
3. Select a language.
4. Customize the template content. (Default fields are automatically filled when a message template is created.)

 **NOTE**

- You can create up to 100 message templates. If this number has been reached, delete unnecessary templates.
- There are two default message templates. If you do not customize any message template, notifications will be sent based on default templates. The default templates cannot be deleted or edited.
- In addition to default fields, the message template also allows you to add custom fields. You need to define values for custom fields for event alarm reporting. For details about the parameters, see the alarm reporting structs in the following message template.
- Custom fields support the JSONPath format. Example: **\$event.metadata.case1** or **\$event.metadata.case[0]**.
- In the upper right corner of the **Body** area, click **Add Variables** to add required variables.
- If you select **Email**, you can click **Preview** to view the final effect. On the **Preview** page, change the message topic if necessary.

Table 6-11 Variables in the default message template

Variable	Description	Definition
Notification Type	Type selected when a notification rule is created, which can be Alarm or Event .	<code>\${event_type}</code>
Severity	Alarm or event severity, which can be Critical , Major , Minor , or Warning .	<code>\${event_severity}</code>
Name	Name of the alarm or event that triggers the notification rule.	<code>\$event.metadata.event_name</code>
Occurred	Time when the alarm or event is triggered.	<code>\${starts_at}</code>
Source	Name of the service corresponding to the alarm or event that triggers the notification rule.	<code>\$event.metadata.resource_provider</code>
Resource Type	Type of the resource selected when you customize a threshold rule or define alarm reporting.	<code>\$event.metadata.resource_type</code>
Resource Identifier	Resource that triggers the alarm or event.	<code>\${resources}</code>
Custom tag	Extended tag.	<code>\$event.metadata.key1</code>
Possible Cause	Cause of the alarm. For non-custom reporting, "NA" is displayed.	<code>\${alarm_probableCause_zh}</code>

Variable	Description	Definition
Additional Info	Additional alarm description, such as the metric name and alarm rule status change.	\${message}
Suggestion	Suggestion on how to handle the alarm. For non-custom reporting, "NA" is displayed.	\${alarm_fix_suggestion_zh}
Custom annotation	Extended annotation.	\$event.annotations.key2

Alarm reporting structs corresponding to the message template

```
{
  "event": {
    "starts_at": 1579420868000,    //${starts_at}
    "ends_at": 1579420868000,
    "timeout": 60000,
    "resource_group_id": "5680587ab6*****755c543c1f",
    "metadata": {
      "event_name": "test",        //${metadata.event_name}
      "event_severity": "Major",  //${metadata.event_severity}
      "event_type": "alarm",      //${metadata.event_type}
      "resource_provider": "ecs", //${metadata.resource_provider}
      "resource_type": "vm",      //${metadata.resource_type}
      "resource_id": "ecs123",
      "key1": "Custom field"      //${event.metadata.key1}
    },
    "annotations": {
      "alarm_probableCause_zh_cn": "possible cause", //${annotations.alarm_probableCause_zh}
      "alarm_fix_suggestion_zh_cn": "fix suggestion", //${annotations.alarm_fix_suggestion_zh}
      "key2": "Custom field"      //${event.annotations.key2}
    }
  }
}
```

5. Click **Confirm**. The message template is created.


----End

More Operations

After creating a message template, you can perform the operations listed in [Table 6-12](#).

Table 6-12 Related operations

Operation	Description
Editing a message template	Click Edit in the Operation column.
Copying a message template	Click Copy in the Operation column.

Operation	Description
Deleting a message template	<ul style="list-style-type: none">To delete a single message template, click Delete in the Operation column in the row that contains the template, and then click Yes on the displayed page.To delete one or more message templates, select them, and click Delete above the template list, and then click Yes on the displayed page. <p>NOTE Before deleting a message template, you need to delete the alarm action rules bound to it.</p>
Searching for a message template	Enter a template name in the search box in the upper right corner and click  .

6.5 Alarm Noise Reduction

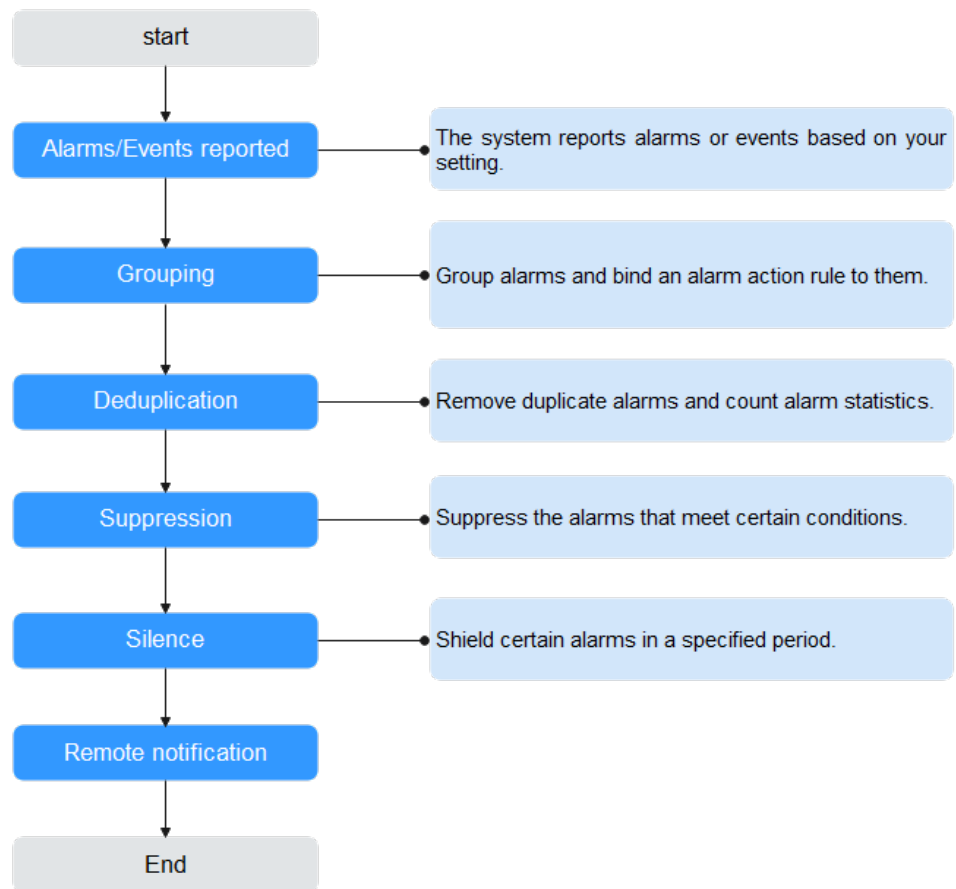
6.5.1 Overview

AOM supports alarm noise reduction. Alarms can be processed based on the alarm noise reduction rules to prevent notification storms.

Alarm noise reduction consists of four parts: grouping, deduplication, suppression, and silence.

AOM uses built-in deduplication rules. The service backend automatically deduplicates alarms. You do not need to manually create rules.

Figure 6-1 Alarm noise reduction process



You need to manually create grouping, suppression, and silence rules. For details, see the following description.

 NOTE

1. This module is used only for message notification. All triggered alarms and events can be viewed on the **Alarm List** and **Event List** pages.
2. All conditions of alarm noise reduction rules are obtained from **metadata** in alarm structs. You can use the default fields or customize your own fields.

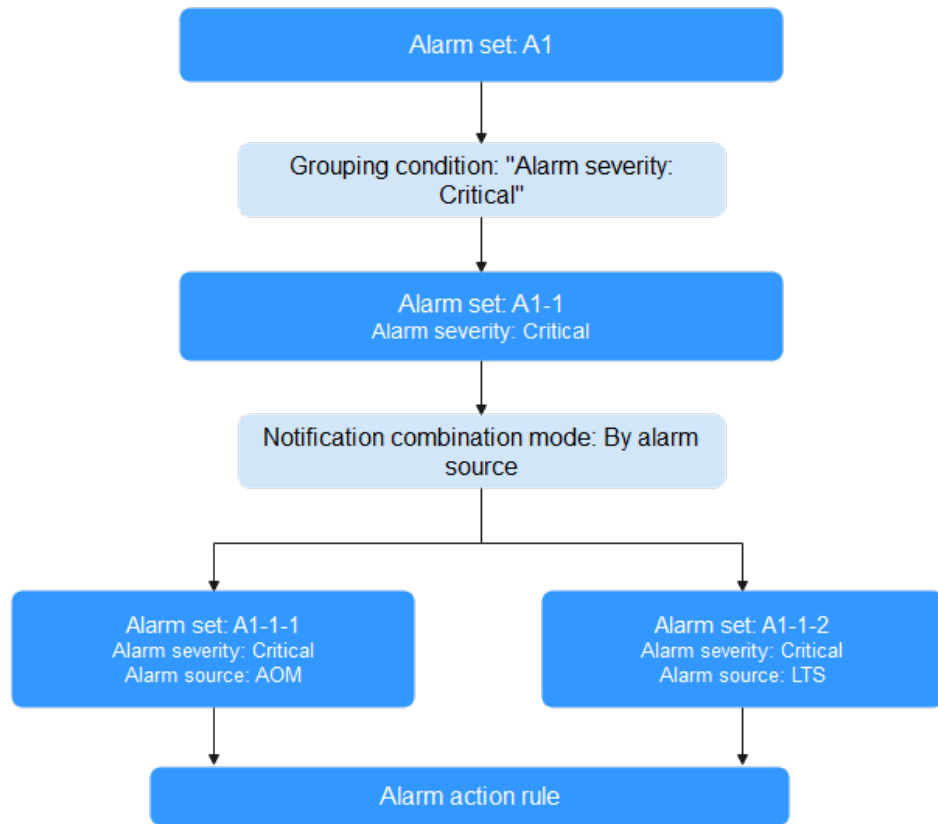
```
{
  "starts_at" : 1579420868000,
  "ends_at" : 1579420868000,
  "timeout" : 60000,
  "resource_group_id" : "5680587ab6*****755c543c1f",
  "metadata" : {
    "event_name" : "test",
    "event_severity" : "Major",
    "event_type" : "alarm",
    "resource_provider" : "ecs",
    "resource_type" : "vm",
    "resource_id" : "ecs123",
    "key1" : "value1" // Alarm tag configured when the alarm rule is created
  },
  "annotations" : {
    "alarm_probableCause_en_us": " Possible causes",
    "alarm_fix_suggestion_en_us": "Handling suggestion"
  }
}
```

6.5.2 Creating a Grouping Rule

You can filter alarm subsets and then group them based on the grouping conditions. Alarms in the same group are aggregated to trigger one notification.

As shown in [Figure 6-2](#), when **Alarm Severity** under **Grouping Condition** is set to **Critical**, the system filters out the critical alarms, and then combines these alarms based on the specified mode. The combined alarms can then be associated with an action rule for sending notifications.

Figure 6-2 Grouping process



Creating a Grouping Rule

You can create up to 100 grouping rules.

Step 1 In the navigation pane, choose **Alarm Center > Alarm Noise Reduction**.

Step 2 On the **Grouping Rules** tab page, click **Create** and set parameters such as the rule name and grouping condition. For details, see [Table 6-13](#).

Table 6-13 Setting a grouping rule

Category	Parameter	Description
-	Rule Name	Grouping rule name, which can contain up to 100 characters and cannot start or end with an underscore (_). Only letters, digits, and underscores are allowed.
	Description	Grouping rule description, which can contain up to 1024 characters.

Category	Parameter	Description
Grouping rules	Grouping Condition	<p>Conditions set to filter alarms. After alarms are filtered out, you can set alarm action rules for them.</p> <p>You can create a maximum of 10 parallel conditions, each of which can contain up to 10 serial conditions. One or more alarm action rules can be set for each parallel condition.</p> <p>Serial conditions are in the AND relationship whereas parallel conditions are in the OR relationship. An alarm must meet all serial conditions under one of the parallel conditions.</p> <p>For example, if two serial conditions (that is, Alarm Severity = Critical and Provider = AOM) are set under a parallel condition, critical AOM alarms are filtered out, and notification actions are performed based on the alarm action rule you set.</p>
Combination rules	Combine Notifications	<p>Notifications for alarms with certain fields to be the same will be combined.</p> <p>Notifications can be combined:</p> <ul style="list-style-type: none"> • By alarm source • By alarm source + severity • By alarm source + all tags
	Initial Wait Time	<p>Interval for sending an alarm notification after alarms are combined for the first time. It is recommended that the time be set to seconds to prevent alarm storms.</p> <p>Value range: 0s to 10 minutes. Recommended: 15s.</p>
	Batch Processing Interval	<p>Waiting time for sending an alarm notification after the combined alarm data changes. It is recommended that the time be set to minutes. If you want to receive alarm notifications as soon as possible, set the time to seconds.</p> <p>The change here refers to a new alarm or an alarm status change.</p> <p>Value range: 5s to 30 minutes. Recommended: 60s.</p>
	Repeat Interval	<p>Waiting time for sending an alarm notification after the combined alarm data becomes duplicate. It is recommended that the time be set to hours.</p> <p>Duplication means that no new alarm is generated and no alarm status is changed while other attributes (such as titles and content) are changed.</p> <p>Value range: 0 minutes to 15 days. Recommended: 1 hour.</p>


Step 3 Click **Create Now**.

----End

More Operations

After creating a grouping rule, perform the operations listed in [Table 6-14](#) if needed.

Table 6-14 Related operations

Operation	Description
Modifying a grouping rule	Click Modify in the Operation column.
Deleting a grouping rule	<ul style="list-style-type: none">To delete a single rule, click Delete in the Operation column in the row that contains the rule.To delete one or more rules, select them and click Delete above the rule list.
Searching for a grouping rule	Enter a rule name in the search box in the upper right corner and click  .

6.5.3 Creating a Suppression Rule

By using suppression rules, you can suppress or block notifications related to specific alarms. For example, when a major alarm is generated, less severe alarms can be suppressed. Another example, when a node is faulty, all other alarms of the processes or containers on this node can be suppressed.

Precautions

If the source alarm corresponding to the suppression condition is cleared before the alarm notification is sent, the suppression rule becomes invalid. For the suppressed object (alarm suppressed by the source alarm), the alarm notification can still be sent as usual.

You can create up to 100 suppression rules.

Creating a Suppression Rule

Step 1 In the navigation pane, choose **Alarm Center > Alarm Noise Reduction**.

Step 2 On the **Suppression Rules** tab page, click **Create** and set parameters such as the rule name and root alarm.

Table 6-15 Setting a suppression rule

Category	Parameter	Description
-	Rule Name	Suppression rule name, which can contain up to 100 characters and cannot start or end with an underscore (_). Only letters, digits, and underscores are allowed.
	Description	Suppression rule description, which can contain up to 1024 characters.
Suppression rules	Source Alarm	Alarm that triggers suppression. You can create up to 10 parallel conditions under Source Alarm , and up to 10 serial conditions under each parallel condition. Serial conditions are in the AND relationship whereas parallel conditions are in the OR relationship. An alarm must meet all serial conditions under one of the parallel conditions. For a serial condition, if Alarm Severity is set to Critical , critical alarms are filtered out as the root alarms.
	Suppressed Alarm	Alarm that is suppressed by the root alarm. Set parameters for the suppressed alarm in the same way that you set parameters for the source alarm. If Serial Condition of Source Alarm is set to Critical and that of Suppressed Alarm is set to Warning , warnings will be suppressed when critical alarms are generated.

Step 3 After you finish setting the parameters, click **Create Now**.

After a suppression rule is created, it will take effect for all alarms that are grouped.


----End

More Operations

After creating a suppression rule, perform the operations listed in [Table 6-16](#) if needed.

Table 6-16 Related operations

Operation	Description
Modifying a suppression rule	Click Modify in the Operation column.

Operation	Description
Deleting a suppression rule	<ul style="list-style-type: none"> To delete a single rule, click Delete in the Operation column in the row that contains the rule. To delete one or more rules, select them and click Delete above the rule list.
Searching for a suppression rule	Enter a rule name in the search box in the upper right corner and click  .

6.5.4 Creating a Silence Rule

You can shield alarm notifications in a specified period. A silence rule takes effect immediately after it is created.

Creating a Silence Rule

You can create up to 100 silence rules.

- Step 1** In the navigation pane, choose **Alarm Center > Alarm Noise Reduction**.
- Step 2** On the **Silence Rules** tab page, click **Create** and set parameters such as the rule name and silence condition.

Table 6-17 Setting a silence rule

Category	Parameter	Description
-	Rule Name	Silence rule name, which can contain up to 100 characters and cannot start or end with an underscore (_). Only letters, digits, and underscores are allowed.
	Description	Silence rule description, which can contain up to 1024 characters.
Silence rules	Silence Condition	<p>Any alarm notifications that meet the silence condition will be shielded.</p> <p>You can create up to 10 parallel conditions under Silence Condition, and up to 10 serial conditions under each parallel condition. Serial conditions are in the AND relationship whereas parallel conditions are in the OR relationship. An alarm must meet all serial conditions under one of the parallel conditions.</p> <p>For a serial condition, if Alarm Severity is set to Critical, critical alarms are shielded.</p>

Category	Parameter	Description
	Silence Time	Time when alarm notifications are shielded. There are two options: <ul style="list-style-type: none"> • Fixed time: Alarm notifications are shielded only in a specified period. • Cycle time: Alarm notifications are shielded periodically.
	Time Zone/ Language	Time zone and language for which alarm notifications are shielded. The time zone and language configured in Preferences are selected by default. You can change them as required.


Step 3 Click **Create Now**.

----End

More Operations

After creating a silence rule, perform the operations listed in [Table 6-18](#) if needed.

Table 6-18 Related operations

Operation	Description
Modifying a silence rule	Click Modify in the Operation column.
Deleting a silence rule	<ul style="list-style-type: none"> • To delete a single rule, click Delete in the Operation column in the row that contains the rule. • To delete one or more rules, select them and click Delete above the rule list.
Searching for a silence rule	Enter a rule name in the search box in the upper right corner and click  .

7 Resource Monitoring

7.1 Application Monitoring

An application is a group of identical or similar components divided based on service requirements. Applications are categorized into system applications and custom applications. The former are discovered based on built-in rules while the latter are discovered based on custom rules.

After application discovery rules are set, AOM automatically discovers applications that meet the rules and monitors related metrics. For details, see [4.2 Configuring Application Discovery Rules](#).

Procedure

Step 1 In the navigation pane, choose **Monitoring > Application Monitoring**.

 **NOTE**

Set filter criteria above the application list to filter applications.

Step 2 Click an application. On the details page that is displayed, manage and monitor components of the application in batches.

You can also view the component list, host list, and alarm analysis result of the current application.

 **NOTE**



In the upper right corner of the **Application Details** page, you can set a time range to query the component, host, or alarm information of the application. If no data exists within the time range, AOM automatically switches to the **Application Monitoring** page.

Step 3 During routine O&M, you can monitor various metrics of applications on the **View Monitor Graphs** tab page.

- **Creating a view template**

AOM provides a default view template (**Application Template**) which can be modified. You can also click **View Template** to customize one.

- **Adding a metric graph**

- You can click  to add a line graph or  to add a digit graph to the view template. You can also delete, move, and copy metric graphs in the view template. For details, see [5.2 Dashboard](#).

- **Adding to a dashboard**

On the application details page, click the **View Monitor Graphs** tab, and choose **More > Add to Dashboard** in the upper right corner to add the view template to the dashboard for monitoring.

Step 4 Perform the following operations if needed:

- **Adding an application**

For identical or similar components that are discovered by default discovery rules or that are not installed with Application Performance Management (APM) probes, you can group them logically, that is, add them to the same application for monitoring.

In the upper right corner of the **Application Monitoring** page, click **Create Application**. On the displayed page, add a custom application discovery rule. For details, see [4.2 Configuring Application Discovery Rules](#). You can monitor the application after adding it. AOM can display O&M information by component. For details, see [7.2 Component Monitoring](#).

----End


7.2 Component Monitoring

Components refer to the services that you deploy, including containers and common processes. For example, a workload on the Cloud Container Engine (CCE) is a component, and the Tomcat running on the VM is also a component.

The component list displays the type, CPU usage, memory usage, and alarm status of each component, helping you learn their running status. You can click a component name to learn more information about the component. AOM supports drill-down from a component to an instance, and then to a container. By viewing the status of each layer, you can implement dimensional monitoring for components.

Step 1 In the navigation pane, choose **Monitoring > Component Monitoring**.

- The component list displays information such as **Component Name**, **Status**, **Application**, **Deployment Mode**, and **Application Discovery Rules**.

- Click  in the upper right corner and select **Hide system component**.
- Set filter criteria above the component list to filter components.

Step 2 Perform the following operations as required:

- **Adding an alias**

If a component name is complex and difficult to identify, you can add an alias for the component.

Click **Add alias** in the **Operation** column to add an alias.

- **Adding a tag**


Tags are identifiers of components. You can distinguish system components from non-system ones based on tags. By default, AOM adds the **System Service** tag to system components (including icagent, css-defender, nvidia-driver-installer, nvidia-gpu-device-plugin, kube-dns, org.tanukisoftware.wrapper.WrapperSimpleApp, evs-driver, obs-driver, sfs-

driver, icwatchdog, and sh). You can click  in the upper right corner to select or deselect **Hide system component**. AOM also allows you to customize tags for easier component management.

In the component list, click **Add tags** in the **Operation** column of the

component, enter a tag, and click  and **OK** to add a tag. You can also mark the component as a system component.

 **NOTE**

- The **Tags** column of the component list is hidden by default. You can click  in the upper right corner and select or deselect **Tags** to show or hide tags.
- **Application Discovery Rules:**
 - **Sys_Rule:** AOM automatically discovers components based on the built-in application discovery rule named **Sys_Rule**. For details, see [Built-in Discovery Rules](#).
 - **Default_Rule:** AOM automatically discovers components based on the built-in application discovery rule named **Default_Rule**. For details, see [Built-in Discovery Rules](#).
 - Custom rules: Their names are customized and not fixed. Applications are discovered based on custom rules.

Step 3 Set filter criteria to search for the desired component.

 **NOTE**

Components cannot be searched by alias.

Step 4 Click the component name. The **Component Details** page is displayed.

 **NOTE**



In the upper right corner of the **Component Details** page, you can set a time range to query the instance, host, or alarm information of the component. If no data exists within the time range, AOM automatically switches to the **Component Monitoring** page.

- On the **Instance List** tab page, view the instance details.

 **NOTE**

Click an instance name to monitor the resource usage and health status.

- On the **Host List** tab page, view the host details.
- On the **Alarm Analysis** tab page, view the alarm details.
- Click the **View Monitor Graphs** tab to monitor the metrics of the component.
 - AOM provides a default view template (**Service Template**) which can be modified. You can also click View Template to customize one.

- You can click  to add a line graph or  to add a digit graph to the view template. You can also delete, move, and copy metric graphs in the view template. For details, see [5.2 Dashboard](#).
 - **Adding to a dashboard**

On the component details page, click the **View Monitor Graphs** tab, and choose **More > Add to Dashboard** in the upper right corner to add the view template to the dashboard for monitoring.
- End

7.3 Host Monitoring

Hosts include the Elastic Cloud Server (ECS) and Bare Metal Server (BMS). AOM monitors the hosts purchased during Cloud Container Engine (CCE) or ServiceStage cluster creation and those directly purchased. Ensure that hosts meet operating system (OS) and version requirements, and the ICAgent is installed on them according to [9.1.1 Installing an ICAgent](#). Otherwise, these hosts cannot be monitored by AOM. In addition, the hosts support both IPv4 and IPv6 addresses.

AOM monitors common system devices such as disks and file systems, and resource usage and health status of hosts and service processes or instances running on them.


Precautions

- A maximum of five tags can be added to a host, and each tag must be unique.
- The same tag can be added to different hosts.
- For hosts created on the CCE or ServiceStage console, you cannot select clusters or create aliases for them.
- The host status can be **Normal**, **Abnormal**, **Warning**, **Silent**, or **Deleted**. The running status of a host is displayed as **Abnormal** when the host is faulty due to network failures, and power off or shut down of the host, or a threshold alarm is reported on the host.

Procedure

Step 1 In the navigation pane, choose **Monitoring > Host Monitoring**.

To view the host list more easily, you can:

- Click  in the upper right corner and select **Hide master host**.
- Set filter criteria above the host list to filter hosts.

Step 2 Perform the following operations as required:


- **Adding an alias**


If a host name is too complex, you can add a simple alias.
In the host list, click **Add alias** in the **Operation** column.

- **Adding a tag**

A tag is the identifier of a host. You can manage and classify hosts by tag. After a tag is added, you can quickly identify, select, or search for a host.

In the host list, choose **More > Add tags** in the **Operation** column, enter a

tag, and click  and **OK** to add a tag. The **Tags** column of the host list is

hidden by default. You can click  in the upper right corner and select or deselect **Tags** to show or hide tags.

- **Synchronizing host data**

In the host list, locate the target host and choose **More > Sync Host Data** in the **Operation** column to synchronize host data.

Step 3 Set filter criteria to search for the desired host.

 **NOTE**

Hosts cannot be searched by alias.

Step 4 Click the host name to enter the **Host Details** page. In the instance list, monitor the resource usage and health status of instances. In addition, click the **View Monitor Graphs** tab to monitor the metrics of the host.



 **NOTE**

In the upper right corner of the **Host Details** page, you can set the time range to query the instance, GPU, NIC, and alarm information of the host. If no data exists within the time range, AOM automatically switches to the **Host Monitoring** page.

- **Creating a view template**

AOM provides a default view template (**Host Template**) which can be modified. You can also click **View Template** to customize one.

- **Adding a metric graph**

– You can click  to add a line graph or  to add a digit graph to the view template. You can also delete, move, and copy metric graphs in the view template. For details, see [5.2 Dashboard](#).

- **Adding to a dashboard**

On the host details page, click the **View Monitor Graphs** tab, and choose **More > Add to Dashboard** in the upper right corner to add the view template to the dashboard for monitoring.

Step 5 Monitor common system devices such as the GPU and NIC of the host.

- Click the **Instance List** tab to view the basic information such as the instance status and type. Click an instance to view its metrics on the details page.
- Click the **GPUs** tab to view the basic information about the GPU of the host. Click a GPU to monitor its metrics on the **View Monitor Graphs** page.
- Click the **NIC** tab to view the basic information about the NIC of the host. Click a NIC to monitor its metrics on the **View Monitor Graphs** page.
- Click the **Disks** tab to view the basic information about the disk of the host. Click a disk to monitor its metrics on the **View Monitor Graphs** page.

- Click the **File System** tab to view the basic information about the file system of the host. Click a disk file partition to monitor its metrics on the **View Monitor Graphs** page.
- Click the **Alarm Analysis** tab to view the alarm details.
- Click the **Disk Partition** tab to view the disk partition type, size, and usage.

 **NOTE**

Disk partitions are supported by CentOS 7.x and EulerOS 2.5.

----End

7.4 Container Monitoring

Container and component monitoring differs in their monitored objects. For component monitoring, workloads deployed using Cloud Container Engine (CCE), applications created using ServiceStage, and components deployed on Elastic Cloud Server (ECS) or Bare Metal Server (BMS) are monitored. For container monitoring, only workloads deployed using CCE and applications created using ServiceStage are monitored. For details, see [7.2 Component Monitoring](#).

7.5 Metric Monitoring

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

Procedure

- Step 1** In the navigation pane, choose **Monitoring > Metric Monitoring**.
- Step 2** Then, click **Add Metric** and select **Dimension** or **Resource** for **Add By**. Select up to 12 metrics to monitor.
- Step 3** Set metric parameters according to [Table 7-1](#), view the metric graph in the upper part of the page, and analyze metric data from multiple dimensions.

Table 7-1 Metric parameters

Parameter	Description
Statistical Mode	Method used to measure metrics. Options: Average, Minimum, Maximum, Sum, and SampleCount . NOTE The number of samples equals to the count of data points.
Statistical Cycle	Interval at which metric data is collected. The statistical cycles that are available for you to select vary according to the time range.





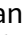

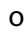
Parameter	Description
Time Range	Time range in which metric data is collected. Options: Last 30 minutes, Last 1 hour, Last 6 hours, Last 1 day, Last 1 week, and Custom.
Refresh Frequency	Interval at which the metric data is refreshed. Options: Refresh manually, 30 seconds auto refresh, 1 minute auto refresh, and 5 minutes auto refresh.

----End

More Operations

You can also perform the operations listed in [Table 7-2](#).

Table 7-2 Related operations

Operation	Description
Hiding/ Showing metric data	After selecting a metric, click  in the Operation column to hide the metric data in the current graph. To show the metric data again, click  in the Operation column.  and  indicate the statuses of metric data.
Adding an alarm rule for a metric	After selecting a metric, click  in the Operation column to create an alarm rule for it.
Copying metric data	After selecting a metric, click  in the Operation column to copy the metric data.
Deleting one or more metrics	<ul style="list-style-type: none"> To delete a metric, click  in the Operation column. To delete one or more metrics, select them and click Delete above the metric list.
Exporting a monitoring report	Click Export Report to export a metric graph as a CSV file to your local PC.

8 Log Management

8.1 Searching for Logs

AOM enables you to quickly query logs, and locate faults based on log sources and contexts.


Step 1 In the navigation pane, choose **Log > Log Search**.




Step 2 On the **Log Search** page, click the **Component**, **System**, or **Host** tab and set filter criteria as prompted.


NOTE

1. You can search for logs by component, system, or host.
 - For component logs, you can set filter criteria such as **Cluster**, **Namespace**, and **Component**. You can also click **Advanced Search** and set filter criteria such as **Instance**, **Host**, and **File Name**, and choose whether to enable **Hide System Component**.
 - For system logs, you can set filter criteria such as **Cluster** and **Host**.
 - For host logs, you can set filter criteria such as **Cluster** and **Host**.
2. Enter a keyword in the search box. Rules are as follows:
 - Enter a keyword between two adjacent delimiters for exact search. By [configuring delimiters](#), you can divide the log content into multiple words and then enter these words to search for logs. If you are not sure whether there are adjacent delimiters, enter a keyword for fuzzy search.
 - Enter a keyword with a question mark (?) or an asterisk (*) for fuzzy match. Do not start a keyword with a question mark or an asterisk. For example, you can enter **ER?OR** or **ER*R**.
 - Enter search criteria containing search operator AND (&&) or OR (||). For example, enter **query logs&&erro*** or **query logs||error**.

Step 3 View the search result of logs.

The search results are sorted based on the log collection time, and keywords in them are highlighted. You can click  in the **Time** column to switch the sorting order.

 indicates the default order.  indicates the ascending order by time (that is, the latest log is displayed at the end).  indicates the descending order by time (that is, the latest log is displayed at the top).

1. Click  on the left of the log list to view details.
2. AOM allows you to view the previous or next logs of a specified log by clicking **View Context** in the **Operation** column, facilitating fault locating. Therefore, you do not need to search for logs in raw files.
 - In the **Display Rows** drop-down list, set the number of rows that display raw context data of the log.



 **NOTE**

For example, select **200** from the **Display Rows** drop-down list.

- If there are 100 logs or more printed prior to a log and 99 or more logs printed following the log, the preceding 100 logs and following 99 logs are displayed as the context.
 - If there are fewer than 100 logs (for example, 90) printed prior to a log and fewer than 99 logs (for example, 80) printed following the log, the preceding 90 logs and following 80 logs are displayed as the context.
- Click **Export Current Page** to export displayed raw context data of the log to a local PC.

 **NOTE**

To ensure that tenant hosts and services run properly, some components (for example, kube-dns) provided by the system will run on the tenant hosts. The logs of these components are also queried during tenant log query.

Step 4 (Optional) Click   in the upper right corner on the **Log Search** page, select the file format, and export the search result to the local PC.

Logs are sorted according to the order set in [Step 3](#) and a maximum of 5000 logs can be exported. For example, when 6000 logs in the search result are sorted in descending order, only the first 5000 logs can be exported.

Logs can be exported in CSV or TXT format. You can select a format as required. If you select the CSV format, detailed information (such as log content, host IP address, and source) can be exported. If you select the TXT format, only log content can be exported. Each row represents a log. If a log contains a large amount of content, you are advised to check the log using a text editor.

----End


8.2 Viewing Log Files

You can quickly view log files of component instances to locate faults.

Viewing Log Files

- Step 1** In the navigation pane, choose **Log > Log Files**.
- Step 2** On the page that is displayed, click the **Component** or **Host** tab and click a component or host name. Information such as the log file name and latest written time is displayed in the log file list on the right.
- Step 3** Click **View** in the **Operation** column of the desired instance. [Table 8-1](#) describes how to check log file details.

Table 8-1 Operations

Operation	Setup	Description
Setting a time range	Date	Click  to select a date.
	Time range	Click the desired time on the time axis to set a time range. You can select only one unit (5 minutes) each time.
Viewing log files	Clear	Click Clear to clear the logs displayed on the screen. Logs displayed on the screen will be cleared, but will not be deleted.
	Viewing real-time logs	<p>The real-time monitoring function is disabled by default. You can click Enable Real-Time Viewing as required. After this function is enabled, the latest written logs can be viewed.</p> <p>The exception in the log records the exceptions that occur during code running. When using logs to locate faults, pay attention to the exception. For real-time log viewing, AOM automatically highlights exception keywords in logs, facilitating fault locating. Such keywords are case-sensitive. For example, exception and Exception are highlighted, but keywords such as EXCEPTION, exCeption, and EXception are not highlighted.</p>

----End

9 Configuration Management

9.1 ICAgent Management

9.1.1 Installing an ICAgent

ICAgents collect metrics, logs, and application performance data in real time. For hosts purchased from the ECS or BMS console, you need to manually install the ICAgent. For hosts purchased from the CCE console, the ICAgent is automatically installed.

 **NOTE**

AOM and LTS use the same ICAgent functions. All metric data collected by ICAgents will be reported to AOM for analysis and processing. However, for logs, only those matching the latest log collection configuration in the system will be collected.

For example, if you configure log collection paths in AOM for ECSs, the previous LTS collection configurations of all ECSs under the resource set become invalid.

The following table describes the ICAgent status.

Table 9-1 ICAgent status

Status	Description
Running	The ICAgent is running properly.
Uninstalled	The ICAgent is not installed. For details about how to install an ICAgent, see 9.1.1 Installing an ICAgent .
Installing	The ICAgent is being installed. This operation takes about 1 minute to complete.
Installation failed	Failed to install the ICAgent. Uninstall the ICAgent according to Uninstalling the ICAgent by Logging In to the Server and then install it again.
Upgrading	The ICAgent is being upgraded. This operation takes about 1 minute to complete.

Status	Description
Upgrade failed	Failed to upgrade the ICAgent. Uninstall the ICAgent according to Uninstalling the ICAgent by Logging In to the Server and then install it again.
Offline	The ICAgent is abnormal due to network problems. Check and restore the network.
Abnormal	The ICAgent is abnormal. Contact technical support.

Prerequisites

Before installing an ICAgent, ensure that the time and time zone of the local browser are consistent with those of the server. If multiple servers are deployed, ensure that the local browser and multiple servers use the same time zone and time. Otherwise, metric data of applications and servers displayed on the UI may be incorrect.

Installation Methods

There are two methods to install an ICAgent. Note that the two methods are not applicable to container nodes created through ServiceStage or CCE. For container nodes, you do not need to manually install an ICAgent. Instead, you only need to perform certain operations when creating clusters or deploying applications.

For details, see [Table 9-2](#).

Table 9-2 Installation methods

Method	Scenario
Initial installation	This method is used when the following condition is met: An ICAgent has never been installed on your server.
Inherited installation	This method is used when the following conditions are met: You need to install ICAgents on multiple servers. An ICAgent has been installed on one of the servers. All the servers are in the same VPC. If the servers are not in the same VPC, bind EIPs to them before using this installation method.

Initial Installation

After you apply for a server and install an ICAgent for the first time, perform the following operations:

- Step 1** Obtain an Access Key ID/Secret Access Key (AK/SK).
- If you have obtained the AK/SK, skip this step.
 - If you have not obtained an AK/SK, [obtain them first](#).

- Step 2** In the navigation pane, choose **Configuration Management > Agent Management**.
- Step 3** Select **Other: custom hosts**, and click **Install ICAgent**.
- Step 4** Click **Copy Command** to copy the installation command.
- Step 5** Use a remote login tool to log in to the target server as the **root** user, and run the following command to disable historical record collection:

```
set +o history
```

- Step 6** Run the copied installation command and enter the obtained AK and SK as prompted.

- Step 7** After the ICAgent is installed, run the following command to enable historical record collection:

```
set -o history
```

 **NOTE**

- If the message **ICAgent install success** is displayed, the ICAgent has been installed in the **/opt/oss/servicemgr/** directory. After the ICAgent has been installed, choose **Configuration Management > Agent Management** to view the ICAgent status.
- If the ICAgent fails to be installed, uninstall the ICAgent according to [Uninstalling the ICAgent by Logging In to the Server](#) and then install it again. If the problem persists, contact technical support.

----End

Inherited Installation

If an ICAgent has been installed on a server and the **ICProbeAgent.zip** installation package exists in the **/opt/ICAgent/** directory of this server, use this method to install an ICAgent on a remote server with a few clicks.

- Step 1** Run the following command (**x.x.x.x** indicates the server IP address) on the server where an ICAgent has been installed:

```
bash /opt/oss/servicemgr/ICAgent/bin/remotelInstall/remote_install.sh -ip  
x.x.x.x
```

- Step 2** Enter the password of the **root** user as prompted.

 **NOTE**

- Inherited installation is not supported when ICAgents are installed using an IAM agency.
- If both the expect tool and the ICAGENT have been installed on the server, an ICAGENT will be installed on the remote server after the preceding command is executed. If an ICAGENT has been installed on the server, but the Expect tool has not, enter the information as prompted for installation.
- Ensure that the **root** user can run the **SSH** or **SCP** command on the server where an ICAGENT has been installed to remotely communicate with the server where an ICAGENT is to be installed.
- If the message **ICAGENT install success** is displayed, the ICAGENT has been installed in the **/opt/oss/servicemgr/** directory. After the ICAGENT has been installed, choose **Configuration Management > Agent Management** to view the ICAGENT status.
- If the ICAGENT fails to be installed, uninstall the ICAGENT according to [Uninstalling the ICAGENT by Logging In to the Server](#) and then install it again. If the problem persists, contact technical support.

----End

Inherited Batch Installation

If an ICAGENT has been installed on a server and the **ICProbeAgent.zip** installation package exists in the **/opt/ICAGENT/** directory of this server, use this method to install ICAGENTS on multiple remote servers with a few clicks.

NOTICE

1. Ensure that you can run the **SSH** and **SCP** commands on the server where an ICAGENT has been installed to communicate with the remote servers where an ICAGENT is to be installed.
2. If you have installed an ICAGENT in a server through an agency, you also need to set an agency for other servers where an ICAGENT is to be installed.
3. Batch installation scripts depend on Python versions. You are advised to implement batch installation on hosts running Python 3.x.
4. Press **Enter** at the end of each line in the **iplist.cfg** file.

Prerequisites

The IP addresses and passwords of all servers on which an ICAGENT is to be installed have been collected, sorted in the **iplist.cfg** file, and uploaded to the **/opt/ICAGENT/** directory on the server where an ICAGENT has been installed. The following is an example of the **iplist.cfg** file, where IP addresses and passwords are separated by spaces.

192.168.0.109 password (Set the password as required.)

192.168.0.39 password (Set the password as required.)

 NOTE

- Because the **iplist.cfg** file contains sensitive information, you are advised to clear the information in time.
- If the passwords of all servers are the same, list IP addresses in the **iplist.cfg** file and enter the password during execution. If the password of an IP address is different from those of other IP addresses, enter the password next to this IP address.
- Batch installation depends on Python 3.x. If the system displays a message indicating that Python cannot be found during the installation, install Python and try again.
- Before the installation, check whether the Python command file exists. If the file does not exist, create a soft link.

Procedure

Step 1 Run the following command on the server where an ICAgent has been installed:

```
bash /opt/oss/servicemgr/ICAgent/bin/remotelInstall/remote_install.sh -  
batchModeConfig /opt/ICAgent/iplist.cfg
```

Enter the preset password of the **root** user as prompted. If the passwords of all IP addresses have been configured in the **iplist.cfg** file, press **Enter** to skip this step. Otherwise, enter the preset password.

```
batch install begin  
Please input default passwd:  
send cmd to 192.168.0.109  
send cmd to 192.168.0.39  
2 tasks running, please wait...  
2 tasks running, please wait...  
2 tasks running, please wait...  
End of install agent: 192.168.0.39  
End of install agent: 192.168.0.109  
All hosts install icagent finish.
```

Wait until the message **All hosts install icagent finish.** is displayed, which indicates that the ICAgent has been installed on all the hosts listed in the configuration file.

Step 2 After the ICAgent has been installed, choose **Configuration Management > Agent Management** to view the ICAgent status.

----End

9.1.2 Upgrading the ICAgent

To ensure better collection experience, AOM will continuously upgrade ICAgent versions. When the system displays a message indicating that a new ICAgent version is available, perform the following operations:

 NOTE

If the ICAgent has a critical bug, the system will upgrade the ICAgent version.

Step 1 In the navigation pane, choose **Configuration Management > Agent Management**.

Step 2 Select **Cluster: xxx** or **Other: custom hosts** from the drop-down list on the right of the page.

Step 3 Upgrade the ICAgent. If you select **Cluster: xxx** in **Step 2**, directly click **Upgrade ICAgent**. In this way, the ICAgent on all hosts in the cluster can be upgraded at

one time. If you select **Other: custom hosts** in **Step 2**, select a desired host and then click **Upgrade ICAgent**.

Step 4 The upgrade takes about 1 minute to complete. When the ICAgent status changes from **Updating** to **Running**, the ICAgent has been upgraded.

 **NOTE**

If the ICAgent state is abnormal after the upgrade or the upgrade fails, log in to the node and run the installation command to reinstall the ICAgent. The overwrite installation mode is supported. Therefore, you can reinstall the ICAgent without uninstallation.

----End

9.1.3 Uninstalling the ICAgent

If the ICAgent on a server is uninstalled, server O&M will be affected, making AOM functions unavailable. Exercise caution when performing this operation.

You can uninstall the ICAgent using either of the following methods:

- **Uninstalling the ICAgent on the AOM Console:** applies to the scenario where the ICAgent has been installed and needs to be uninstalled.
- **Uninstalling the ICAgent by Logging In to the Server:** applies to the scenario where the ICAgent fails to be installed and needs to be uninstalled.
- **Remotely Uninstalling the ICAgent:** applies to the scenario where the ICAgent has been installed and needs to be remotely uninstalled.
- **Uninstalling the ICAgent in Batches:** applies to the scenario where the ICAgent has been installed and needs to be uninstalled in batches.

Uninstalling the ICAgent on the AOM Console

Step 1 In the navigation pane, choose **Configuration Management > Agent Management**.

Step 2 Select **Other: custom hosts** from the drop-down list on the right of the page.

Step 3 Select one or more servers where the ICAgent is to be uninstalled, and click **Uninstall ICAgent**. In the **Uninstall ICAgent** dialog box, click **OK**.

The uninstallation takes about 1 minute to complete. When the ICAgent status changes from **Uninstalling** to **Uninstall**, the ICAgent has been uninstalled.

----End

Uninstalling the ICAgent by Logging In to the Server

Step 1 Log in as the **root** user to the server where the ICAgent is to be uninstalled.

Step 2 Run the following command to uninstall the ICAgent:

```
bash /opt/oss/servicemgr/ICAgent/bin/manual/uninstall.sh;
```

Step 3 If the message **ICAgent uninstall success** is displayed, the ICAgent has been uninstalled.

----End

Remotely Uninstalling the ICAgent

In addition to the preceding method, you can use a method similar to [Inherited Installation](#) to remotely uninstall the ICAgent.

Step 1 Run the following command (*x.x.x.x* indicates the server IP address) on the server where the ICAgent has been installed:

```
bash /opt/oss/servicemgr/ICAgent/bin/remoteUninstall/remote_uninstall.sh -  
ip x.x.x.x
```

Step 2 Enter the password of the **root** user as prompted.

NOTE

- If both the expect tool and the ICAgent have been installed on the server, the ICAgent will be uninstalled from the remote server after the preceding command is executed. If the ICAgent has been installed on the server, but the Expect tool has not, enter the information as prompted for installation.
- Ensure that the **root** user can run the **SSH** or **SCP** command on the server where the ICAgent has been installed to remotely communicate with the server where the ICAgent is to be uninstalled.
- If the message **ICAgent uninstall success** is displayed, the ICAgent has been uninstalled. After the ICAgent has been uninstalled, choose **Configuration Management > Agent Management** to view the ICAgent status.

----End

Uninstalling the ICAgent in Batches

If the ICAgent has been installed on a server and the **ICProbeAgent.zip** installation package exists in the **/opt/ICAgent/** directory of this server, use this method to uninstall the ICAgent from multiple remote servers in batches with a few clicks.

NOTICE

The servers must belong to the same Virtual Private Cloud (VPC) and network segment.

Prerequisites

The IP addresses and passwords of all servers from which the ICAgent is to be uninstalled have been collected, sorted in the **iplist.cfg** file, and uploaded to the **/opt/ICAgent/** directory on the server where the ICAgent has been installed. The following is an example of the **iplist.cfg** file, where IP addresses and passwords are separated by spaces.

```
192.168.0.109 password (Set the password as required.)
```

```
192.168.0.39 password (Set the password as required.)
```

 NOTE

- Because the **iplist.cfg** file contains sensitive information, you are advised to clear the information in time.
- If the passwords of all servers are the same, list IP addresses in the **iplist.cfg** file and enter the password during execution. If the password of an IP address is different from those of other IP addresses, enter the password next to this IP address.
- You need to press **Enter** at the end of each line in the **iplist.cfg** file.

Procedure

Step 1 Run the following command on the server where the ICAgent has been installed:

```
bash /opt/oss/servicemgr/ICAgent/bin/remoteUninstall/remote_uninstall.sh -  
batchModeConfig /opt/ICAgent/iplist.cfg
```

Enter the default password of the **root** user as prompted. If the passwords of all IP addresses have been configured in the **iplist.cfg** file, press **Enter** to skip this step. Otherwise, enter the default password.

```
batch uninstall begin  
Please input default passwd:  
send cmd to 192.168.0.109  
send cmd to 192.168.0.39  
2 tasks running, please wait...  
End of uninstall agent: 192.168.0.109  
End of uninstall agent: 192.168.0.39  
All hosts uninstall icagent finish.
```

Wait until the message **All hosts uninstall icagent finish.** is displayed, which indicates that the ICAgent has been uninstalled from all the hosts listed in the configuration file.

Step 2 After the ICAgent has been uninstalled, choose **Configuration Management > Agent Management** to view the ICAgent status.

----End

9.2 Log Configuration

9.2.1 Setting the Log Quota

Step 1 Log in to the AOM console. In the navigation pane, choose **Configuration Management > Log Configuration**.

Step 2 On the **Quota Management** page, view the log size and retention period.

The default log retention period is 7 days and can be changed as required (maximum: 30 days).

----End

9.2.2 Configuring Delimiters

AOM enables you to divide the log content into multiple words for search by configuring delimiters. By default, AOM provides the following delimiters:

```
, ";=() []{}@<>/:\n\t\r
```


If default delimiters cannot meet requirements, customize delimiters according to the following procedure.

Precautions




Delimiters are applicable only to the logs generated after the time when the delimiters are configured. Earlier logs are processed based on earlier delimiters.

Procedure

Step 1 In the navigation pane, choose **Configuration Management > Log Configuration**, and click the **Delimiter Configuration** tab.

Step 2 Configure delimiters.

You can configure delimiters using the following methods: If you use both methods at the same time, the union set will be selected.

- Customize delimiters. Specifically, click , enter a delimiter in the text box, and click .
- Use ASCII code. Specifically, click **Add Special Delimiters**, enter the ASCII value according to [ASCII Comparison Table](#), and click .

Step 3 Preview the log content.

Enter the log content to be previewed in the text box and click **Preview**.

Step 4 Confirm the configuration and click **OK**.

NOTE

Click **Reset** to restore the default configuration. Default delimiters are as follows:

```
,";=()[]{}@&<>/: \n\t\r
```

----End

ASCII Comparison Table

Table 9-3 ASCII comparison table

ASCII Value	Control Character	ASCII Value	Control Character	ASCII Value	Control Character	ASCII Value	Control Character
0	NUL (Null)	32	Space	64	@	96	`
1	SOH (Start of heading)	33	!	65	A	97	a
2	STX (Start of text)	34	"	66	B	98	b
3	ETX (End of text)	35	#	67	C	99	c

ASCII Value	Control Character	ASCII Value	Control Character	ASCII Value	Control Character	ASCII Value	Control Character
4	EOT (End of transmission)	36	\$	68	D	100	d
5	ENQ (Enquiry)	37	%	69	E	101	e
6	ACK (Acknowledge)	38	&	70	F	102	f
7	BEL (Bell)	39	'	71	G	103	g
8	BS (Backspace)	40	(72	H	104	h
9	HT (Horizontal tab)	41)	73	I	105	i
10	LF (Line feed)	42	*	74	J	106	j
11	VT (Vertical tab)	43	+	75	K	107	k
12	FF (Form feed)	44	,	76	L	108	l
13	CR (Carriage return)	45	-	77	M	109	m
14	SO (Shift out)	46	.	78	N	110	n
15	SI (Shift in)	47	/	79	O	111	o
16	DLE (Data link escape)	48	0	80	P	112	p
17	DC1 (Device control 1)	49	1	81	Q	113	q
18	DC2 (Device control 2)	50	2	82	R	114	r

ASCII Value	Control Character	ASCII Value	Control Character	ASCII Value	Control Character	ASCII Value	Control Character
19	DC3 (Device control 3)	51	3	83	S	115	s
20	DC4 (Device control 4)	52	4	84	T	116	t
21	NAK (Negative acknowledge)	53	5	85	U	117	u
22	SYN (Synchronous suspension)	54	6	86	V	118	v
23	ETB (End of transmission block)	55	7	87	W	119	w
24	CAN (Cancel)	56	8	88	X	120	x
25	EM (End of medium)	57	9	89	Y	121	y
26	SUB (Substitute)	58	:	90	Z	122	z
27	ESC (Escape)	59	;	91	[123	{
28	FS (File separator)	60	<	92	/	124	
29	GS (Group separator)	61	=	93]	125	}
30	RS (Record separator)	62	>	94	^	126	~
31	US (Unit separator)	63	?	95	_	127	DEL (Delete)

9.2.3 Setting Log Collection

You can enable or disable log collection as required to reduce memory, database, and disk space usage.

Configuring Log Collection

Before enabling this function, ensure that you have installed the ICAgent on an Elastic Cloud Server (ECS) according to [9.1.1 Installing an ICAgent](#).

Step 1 Log in to the AOM console. In the navigation pane, choose **Configuration Management > Log Configuration**. Then, click the **Log Switch** tab.

Step 2 Enable or disable log collection.

 **NOTE**

- The log collection function is enabled by default. If you do not need to collect logs, disable this function to reduce resource usage.
- After the log collection function is disabled, ICAgents will stop collecting logs, and this function on the LTS console will also be disabled.

----End

9.3 Quota Configuration

Step 1 Log in to the AOM console.

Step 2 Choose **Configuration Management > Quota Configuration**.

Step 3 Check the metric quota.

Earlier metrics will be deleted when the metric quota is exceeded.

----End

9.4 Metric Configuration

You can determine whether to enable **Metric Collection** to collect metrics (excluding custom metrics).

Before enabling this function, ensure that you have installed the ICAgent on an Elastic Cloud Server (ECS) according to [9.1.1 Installing an ICAgent](#).

Step 1 Log in to the AOM console. In the navigation pane, choose **Configuration Management > Metric Configuration**.

Step 2 Enable or disable metric collection.

 **NOTE**

- After metric collection is disabled, ICAgents will stop collecting metric data and related metric data will not be updated. However, custom metrics can still be reported.

----End

10 Auditing

10.1 Operations Logged by CTS

AOM is a one-stop O&M platform that monitors applications and resources in real time. By analyzing dozens of metrics and correlation between alarms and logs, AOM helps O&M personnel quickly locate faults.

You can use AOM to comprehensively monitor and uniformly manage servers, storage, networks, web containers, and applications hosted in Docker and Kubernetes. This effectively prevents problems and helps O&M personnel locate faults in minutes, reducing O&M costs. Also, AOM provides unified APIs to interconnect in-house monitoring systems or report systems. Unlike traditional monitoring systems, AOM monitors services by applications. It meets enterprises' requirements for high efficiency and fast iteration, provides effective IT support for their services, and protects and optimizes their IT assets, enabling enterprises to achieve strategic goals and maximize value. With CTS, you can record operations associated with AOM for future query, audit, and backtracking.

 **NOTE**

pe traces actually record AOM operations, but these operations are performed through CCE or ServiceStage.

Table 10-1 Operations logged by CTS

Operation	Resource Type	Event Name
Creating a dashboard	ams	add-view-action
Modifying a dashboard	ams	update-view-action
Deleting a dashboard	ams	deleteDashboard
Creating a threshold	ams	addThreshold
Modifying a threshold	ams	updateThreshold
Deleting a threshold	ams	deleteThreshold
Deleting a subscription rule	apminventory	deleteSubscribeRule

Operation	Resource Type	Event Name
Modifying a subscription rule name	apminventory	updateSubscribeName
Creating a subscription rule	apminventory	createSubscribeRule
Enabling the pay-per-use edition	OpenOrClosePro Service	openProBillingService
Disabling the pay-per-use edition	OpenOrClosePro Service	closeProBillingService
Deleting a threshold rule	threshold_rules_v2	deleteOneAlarmById
Deleting threshold rules in batches	threshold_rules_v2	deleteAlarmRules
Modifying a threshold rule	threshold_rules_v2	updateAlarm
Creating a threshold rule	threshold_rules_v2	addAlarmForDT
Modifying an event alarm rule	event2alarm_rule	updateEvent2AlarmRule
Creating an event alarm rule	event2alarm_rule	addEvent2AlarmRule
Deleting an event alarm rule	event2alarm_rule	deleteEvent2AlarmRule
Installing a collector	icmgr	icagentInstall
Upgrading a collector	icmgr	icagentUpgrade
Upgrading a probe	icmgr	pinPointUpgrade
Uninstalling a collector	icmgr	lcagentUninstall
Setting metric and log collection	icmgr	metricAndLogSwitches
Delivering configuration	icmgr	webIcAgentEvent
Clearing an alarm	pushEvents	clearEvents
Creating an alarm action rule	actionRule	addActionRule
Modifying an alarm action rule	actionRule	updateActionRule
Deleting an alarm action rule	actionRule	delActionRule

Operation	Resource Type	Event Name
Creating a message template	notificationTemplate	addNotificationTemplate
Modifying a message template	notificationTemplate	updateTemplate
Deleting a message template	notificationTemplate	delTemplate
Creating a grouping rule	groupRule	addGroupRule
Modifying a grouping rule	groupRule	updateGroupRule
Deleting a grouping rule	groupRule	delGroupRule
Creating a suppression rule	inhibitRule	addInhibitRule
Modifying a suppression rule	inhibitRule	updateInhibitRule
Deleting a suppression rule	inhibitRule	delInhibitRule
Creating a silence rule	muteRule	addMuteRule
Modifying a silence rule	muteRule	updateMuteRule
Deleting a silence rule	muteRule	delMuteRule
Creating or modifying an application discovery rule	apminventory	addOrUpdateAppRules
Deleting an application discovery rule	apminventory	deleteAppRules
Modifying the alias or tag of an application, host, or component	apminventory	updateInventoryTag
Creating a policy group	pe	createPolicyGroup
Deleting a policy group	pe	deletePolicyGroup
Updating a policy group	pe	updatePolicyGroup
Enabling a policy group	pe	enablePolicyGroup
Disabling a policy group	pe	disablePolicyGroup

Operation	Resource Type	Event Name
Creating a policy	pe	createPolicy
Deleting a policy	pe	deletePolicy
Updating a policy	pe	updatePolicy
Enabling a policy	pe	enablePolicy
Disabling a policy	pe	disablePolicy
Updating the aging period	als	updateLogStorageSetting

10.2 Querying Real-Time Traces


Scenarios




After you enable CTS and the management tracker is created, CTS starts recording operations on cloud resources. After a data tracker is created, the system starts recording operations on data in OBS buckets. CTS stores operation records generated in the last seven days.

This section describes how to query and export operation records of the last seven days on the CTS console.


- [Viewing Real-Time Traces in the Trace List of the New Edition](#)
- [Viewing Real-Time Traces in the Trace List of the Old Edition](#)



Viewing Real-Time Traces in the Trace List of the New Edition

1. Log in to the management console.
2. Click  in the upper left corner and choose Management & Deployment > **Cloud Trace Service**. The CTS console is displayed.
3. Choose **Trace List** in the navigation pane on the left.
4. On the **Trace List** page, use advanced search to query traces. You can combine one or more filters.
 - **Trace Name:** Enter a trace name.
 - **Trace ID:** Enter a trace ID.
 - **Resource Name:** Enter a resource name. If the cloud resource involved in the trace does not have a resource name or the corresponding API operation does not involve the resource name parameter, leave this field empty.
 - **Resource ID:** Enter a resource ID. Leave this field empty if the resource has no resource ID or if resource creation failed.
 - **Trace Source:** Select a cloud service name from the drop-down list.
 - **Resource Type:** Select a resource type from the drop-down list.

- **Operator:** Select one or more operators from the drop-down list.
 - **Trace Status:** Select **normal**, **warning**, or **incident**.
 - **normal:** The operation succeeded.
 - **warning:** The operation failed.
 - **incident:** The operation caused a fault that is more serious than the operation failure, for example, causing other faults.
 - Time range: Select **Last 1 hour**, **Last 1 day**, or **Last 1 week**, or specify a custom time range.
5. On the **Trace List** page, you can also export and refresh the trace list, and customize the list display settings.
- Enter any keyword in the search box and press Enter to filter desired traces.
 - Click **Export** to export all traces in the query result as an .xlsx file. The file can contain up to 5000 records.
 - Click  to view the latest information about traces.
 - Click  to customize the information to be displayed in the trace list. If **Auto wrapping** is enabled (), excess text will move down to the next line; otherwise, the text will be truncated. By default, this function is disabled.
6. For details about key fields in the trace structure, see section "Trace References" > "Trace Structure" and section "Trace References" > "Example Traces".
7. (Optional) On the **Trace List** page of the new edition, click **Go to Old Edition** in the upper right corner to switch to the **Trace List** page of the old edition.

Viewing Real-Time Traces in the Trace List of the Old Edition

1. Log in to the management console.
2. Click  in the upper left corner and choose **Management & Deployment** > **Cloud Trace Service**. The CTS console is displayed.
3. Choose **Trace List** in the navigation pane on the left.
4. Each time you log in to the CTS console, the new edition is displayed by default. Click **Go to Old Edition** in the upper right corner to switch to the trace list of the old edition.
5. Set filters to search for your desired traces. The following filters are available:
 - **Trace Type**, **Trace Source**, **Resource Type**, and **Search By**: Select a filter from the drop-down list.
 - If you select **Resource ID** for **Search By**, specify a resource ID.
 - If you select **Trace name** for **Search By**, specify a trace name.
 - If you select **Resource name** for **Search By**, specify a resource name.
 - **Operator**: Select a user.

- **Trace Status:** Select **All trace statuses**, **Normal**, **Warning**, or **Incident**.
 - Time range: You can query traces generated during any time range in the last seven days.
 - Click **Export** to export all traces in the query result as a CSV file. The file can contain up to 5000 records.
6. Click **Query**.
 7. On the **Trace List** page, you can also export and refresh the trace list.
 - Click **Export** to export all traces in the query result as a CSV file. The file can contain up to 5000 records.
 - Click  to view the latest information about traces.
 8. Click  on the left of a trace to expand its details.

Trace Name	Resource Type	Trace Source	Resource ID	Resource Name	Trace Status	Operator	Operation Time	Operation
createDockerConfig	dockerlogincmd	SWR	-	dockerlogincmd	normal		Nov 16, 2023 10:54:04 GMT+08:00	View Trace

```

request
trace_id
code: 200
trace_name: createDockerConfig
resource_type: dockerlogincmd
trace_rating: normal
api_version
message: createDockerConfig, Method: POST Url=/v2/manage/utils/secret, Reason:
source_ip
domain_id
trace_type: ApiCall
            
```

9. Click **View Trace** in the **Operation** column. The trace details are displayed.

View Trace ×

```

{
  "request": "",
  "trace_id": "",
  "code": "200",
  "trace_name": "createDockerConfig",
  "resource_type": "dockerlogincmd",
  "trace_rating": "normal",
  "api_version": "",
  "message": "createDockerConfig, Method: POST Url=/v2/manage/utils/secret. Reason:",
  "source_ip": "",
  "domain_id": "",
  "trace_type": "ApiCall",
  "service_type": "SWR",
  "event_type": "system",
  "project_id": "",
  "response": "",
  "resource_id": "",
  "tracker_name": "system",
  "time": "Nov 16, 2023 10:54:04 GMT+08:00",
  "resource_name": "dockerlogincmd",
  "user": {
    "domain": {
      "name": "",
      "id": ""
    }
  }
}
            
```

10. For details about key fields in the trace structure, see section "Trace References" > "Trace Structure" and section "Trace References" > "Example Traces" in the *CTS User Guide*.
11. (Optional) On the **Trace List** page of the old edition, click **New Edition** in the upper right corner to switch to the **Trace List** page of the new edition.

11 Upgrading to AOM 2.0

This section describes how to migrate data from AOM 1.0 to AOM 2.0. Currently, only alarm rule and collector upgrades are supported.

Introduction

- **Collector Upgrade**
After the upgrade, the process discovery capability is enhanced and the collector can automatically adapt to functions related to CMDB, and monitoring center.
- **Alarm Rule Upgrade**
After the upgrade, alarm rule data is smoothly switched from AOM 1.0 to AOM 2.0, and is automatically adapted to alarm rule functions of AOM 2.0.

Collector Upgrade

- Step 1** Log in to the AOM 1.0 console.
- Step 2** In the navigation pane, choose **Configuration Management > Agent Management**.
- Step 3** Select **Other: custom hosts** from the drop-down list on the right of the page.
- Step 4** Select a host and click **Upgrade ICAgent**.
- Step 5** Select a target version from the drop-down list and click **OK**.
- Step 6** Wait for the upgrade. This process takes about one minute. When the ICAgent status changes from **Updating** to **Running**, the ICAgent has been upgraded.

NOTE

If the ICAgent is abnormal after the upgrade or if the upgrade fails, log in to the host and run the installation command again. There is no need for you to uninstall the original ICAgent.

----End

Alarm Rule Upgrade

Step 1 Log in to the AOM 1.0 console.

Step 2 In the navigation pane on the left, choose **Alarm Center > Alarm Rules**.

Step 3 Select one or more alarm rules and click **Migrate to AOM 2.0** above the rule list.

NOTICE

- Migration cannot be undone. Exercise caution when performing this operation.
- If the alarm rules to be migrated depend on alarm templates, these alarm templates will also be migrated.

Step 4 In the displayed dialog box, click **Confirm**. The selected alarm rules will be migrated to AOM 2.0 in batches.

----End

12 FAQs

12.1 User FAQs

Why Is Monitoring Data Not Displayed in Real Time on the AOM Page After Resources Are Created?

After you create resources such as hosts, applications, components, and processes, the ICAgent reports their monitoring data every 10 minutes. Only when a report period ends, can monitoring data be displayed on the AOM page.

Why Is the Resource Status Displayed as Normal on the AOM Page After Resources Are Deleted?

After you delete a resource such as a host or workload from a Cloud Container Engine (CCE) cluster, the resource status is still **Normal** on the **Host Monitoring** or **Container Monitoring** page of AOM. This is normal. The status of the deleted resource will be changed to **Deleted** 30 minutes later.

What Can I Do If the ICAgent Fails to Be Upgraded?

In the custom cluster scenario, if the ICAgent fails to be upgraded, log in to the VM node and directly run the ICAgent installation command again.

Because the ICAgent supports overwriting installation, directly reinstall the ICAgent without uninstallation.

What Types of Log Files Can Be Collected?

If you specify a directory, all **.log**, **.trace**, and **.out** log files in this directory are collected by default. If you specify a log file, only this file is collected. The specified file must be a text file. Other types of log files, such as binary log files, cannot be collected.

Does the ICAgent Consume Lots of Resources Such as Memory and CPU?

- AOM collects basic metrics, including CPU and memory of VMs, containers, and processes.

Resource consumption: The resource consumption of the ICAgent is related to the number of containers and processes. In normal service traffic, the ICAgent consumes about 30 MB memory and 3% single-core CPU.

Usage restriction: Ensure that the number of containers running on a single node is less than 1000.

Protection mechanism:

- The ICAgent consumes a maximum of two CPU cores.
- When the memory consumed by the ICAgent exceeds $\min\{4\text{ GB, node physical memory}/2\}$, AOM restarts the ICAgent for protection.

 **NOTE**

$\min\{4\text{ GB, node physical memory}/2\}$ indicates the smaller value between 4 GB and $\text{node physical memory}/2$.

- AOM also collects log files, including syslog, standard container output, user configuration path, and container mounting files.

Resource consumption: The resource consumption of the ICAgent is closely related to the log volume, number of files, network bandwidth, and backend service processing capability.

How Does AOM Obtain a Custom Host IP Address on the Agent Management Page?

By default, AOM traverses all NICs on the VM and obtains the IP addresses of the Ethernet, bond, and wireless NICs based on priorities in descending order. To ensure that AOM obtains the IP address of a specific NIC, set the **IC_NET_CARD=Desired NIC name** environment variable when starting the ICAgent.

Example:

1. Add **export IC_NET_CARD=eth2** to **/etc/profile**.
2. Run the **source /etc/profile** command to make the environment variable effective in the shell.
3. Go to the **/opt/oss/servicemgr/ICAgent/bin/manual/** directory, and stop and then restart the ICAgent.

```
bash mstop.sh
```

```
bash mstart.sh
```

4. Check whether the environment variable is properly transferred to the application.

```
strings /proc/{icagentprocid}/environ | grep IC_NET_CARD
```

 **NOTE**

- If the IP address displayed on ICAgent is **127.0.0.1**, the ICAgent may fail to obtain the local IP address during startup. This problem may occur when a VM is powered off and then restarted. To solve the problem, you only need to restart the ICAgent.
- If the IP address of your host changes (for example, a new IP address is allocated during renewal), the original IP address may be displayed on the agent management page. To solve the problem, you only need to restart the ICAgent.

What Can I Do If the ICAgent Fails to Be Installed in the Windows Environment and the "SERVICE STOP" Message Is Displayed?

Symptom: The ICAgent fails to be installed in the Windows environment and the "SERVICE STOP" message is displayed. No ICAgent task exists in the task manager. No ICAgent service exists in the system service list. When the **sc query icagent** command is executed, a message is displayed, indicating that no ICAgent is found.

Cause: Antivirus software, such as 360 Total Security, blocks registration of the ICAgent service.

Solution:

1. Check whether antivirus software, such as 360 Total Security, is running.
2. First close the antivirus software and install the ICAgent again.

 **NOTE**

In the Windows environment, manually add log collection paths. The ICAgent can collect **.log**, **.trace**, and **.out** files, but does not collect binary files or Windows system logs.

12.2 Consultation FAQs

12.2.1 What Are the Usage Restrictions of AOM?

OS Usage Restrictions

AOM supports multiple operating systems (OSs). When creating a host, ensure that its OS meets the requirements in [Table 12-1](#). Otherwise, the host cannot be monitored by AOM.

Table 12-1 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			

OS	Version				
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 12-2](#).

Table 12-2 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	<ul style="list-style-type: none"> • 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	Metric data can be stored in the database for up to 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.

Category	Object	Usage Restrictions
	Dimension	A maximum of 20 dimensions can be configured for a metric.
	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	<ul style="list-style-type: none"> 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.
Log	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.
	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.

Category	Object	Usage Restrictions
	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.
	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	<p>ICAgent uses multiple mechanisms to ensure log collection reliability and prevent data loss. However, logs may be lost in the following scenarios:</p> <ul style="list-style-type: none"> • 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 15 days.
	Event	You can query the events generated in the last 15 days.
-	Application discovery rule	You can create a maximum of 100 application discovery rules.

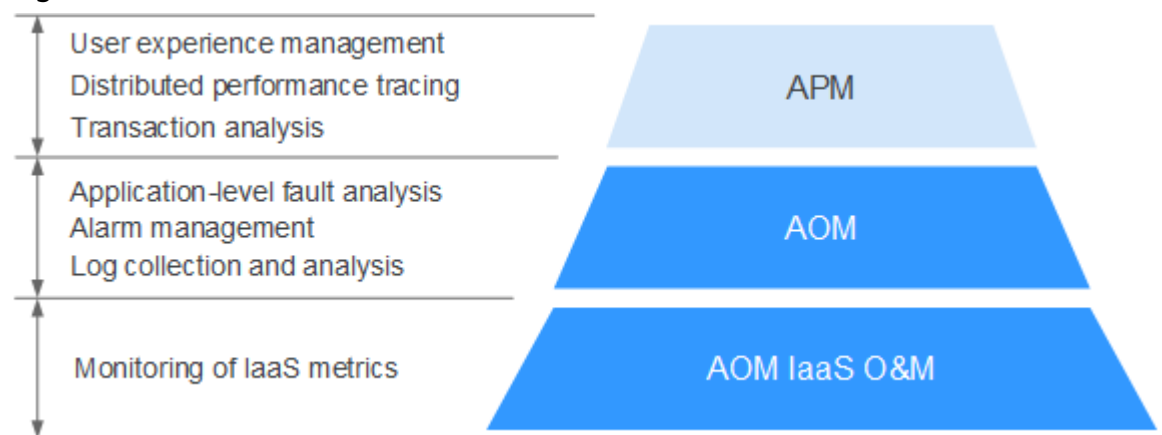
Service Usage Restrictions

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.

12.2.2 What Are the Differences Between AOM and APM?

AOM and Application Performance Management (APM) belong to the multi-dimensional O&M solution and share the ICAgent collector. AOM provides application-level fault analysis, alarm management, and log collection and analysis capabilities, which effectively prevent problems and help O&M personnel quickly locate faults, reducing O&M costs. APM provides user experience management, distributed performance tracing, and transaction analysis capabilities, which help O&M personnel quickly locate and resolve faults and performance bottlenecks in a distributed architecture, optimizing user experience. AOM provides basic O&M capabilities. APM is a supplement to AOM.

Figure 12-1 Multi-dimensional O&M solution



12.2.3 How Do I Distinguish Alarms from Events?

Similarities Between Alarms and Events

Both alarms and events are the information reported to AOM when the status of AOM or an external service (such as ServiceStage or CCE) changes.

Differences Between Alarms and Events

- Alarms are reported when AOM or an external service (such as ServiceStage or CCE) is abnormal or may cause exceptions. Alarms must be handled. Otherwise, service exceptions may occur.
- Events generally carry some important information. They are reported when AOM or an external service (such as ServiceStage or CCE) has some changes. Such changes do not necessarily cause service exceptions. Events do not need to be handled.

12.2.4 What Is the Relationship Between the Time Range and Statistical Cycle?

For Application Operations Management (AOM), a maximum of 1440 data points can be returned for a single metric query. The relationship between the time range and statistical cycle is as follows:

Maximum time range = Statistical cycle x 1440

If you select a time range shorter than or equal to the maximum time range, all the statistical cycles that meet the preceding formula can be selected. For example, if you query metrics in the last 1 hour, the available statistical cycles are 1 minute and 5 minutes.

Table 12-3 shows the relationship between the time range and statistical cycle.

Table 12-3 Relationship between the time range and statistical cycle

Time Range	Statistical Cycle
Last 1 hour	1 minute or 5 minutes
Last 6 hours	1 minute, 5 minutes, or 1 hour
Last 1 day	
Last 7 days	1 hour or 1 day NOTE 1 day is only for the metrics generated based on log statistical rules.
Last 15 days	1 hour or 1 day NOTE 1 day is only for the metrics generated based on log statistical rules.
Last 30 days	
Last 3 months	
Last 6 months	
Last 9 months	
Last 12 months	

12.2.5 Does AOM Display Logs in Real Time?

The logs displayed on Application Operations Management (AOM) are near real-time logs, of which the latency is in seconds.

There is a time interval between log collection and processing. If the number of logs is small, the latency is about 10s. If the number of logs is large, the latency is much longer.

12.2.6 How Can I Do If I Cannot Receive Any Email Notification After Configuring a Threshold Rule?

The notification configuration may be incorrect. You can check whether the alarm notification function is enabled, and whether a topic and an Application Performance Management (APM) policy are selected.

12.2.7 Why Are Connection Channels Required?

Different Virtual Private Clouds (VPCs) cannot communicate with each other. To solve this problem, create an application in the VPC where the data subscription application resides and set it to a VPC endpoint service. Then, create a VPC endpoint in the VPC where Distributed Message Service (DMS) resides. A connection channel can be established between the VPC endpoint and VPC endpoint service for communication.

12.3 Usage FAQs

12.3.1 What Can I Do If I Do Not Have the Permission to Access SMN?

When you log in to AOM and create or modify a threshold rule, notification rule, or static threshold template as an IAM user, the message "Sorry, you do not have the permission to access Simple Message Notification (SMN)" is displayed below **Topic**.

Problem Analysis

- **Cause:** You log in to AOM as an IAM user, but this user does not have the permission to access SMN.
- **Impact:** You cannot receive notifications by email and Short Message Service (SMS) message.

Solution

Contact the administrator (account to which the IAM user belongs) to add the SMN access permission. To add the permission, do as follows:

Log in to IAM as the administrator, and add the SMN access permission to the IAM user.

12.3.2 What Can I Do If Resources Are Not Running Properly?

The resource status includes **Normal**, **Warning**, **Alarm**, and **Silent**. **Warning**, **Alarm**, or **Silent** may result in resource exceptions. You can analyze and rectify exceptions based on the following suggestions.

Warning

If a minor alarm or warning exists, the resource status is **Warning**.

Suggestion: Handle problems based on alarm details.

Alarm

If a critical or major alarm exists, the resource status is **Alarm**.

Suggestion: Handle problems based on alarm details.

Silent

If the ICAgent fails to collect resource metrics, the resource status is **Silent**. The causes include but are not limited to:

- **Cause 1: The ICAgent is abnormal.**

Suggestion: In the navigation pane, choose **Configuration Management > Agent Management**. On the page that is displayed, check the ICAgent status. If the status is not **Running**, the ICAgent is not installed or abnormal. For details about how to solve the problem, see [Table 12-4](#).

Table 12-4 ICAgent troubleshooting suggestions

Status	Suggestion
Uninstalled	Install an ICAgent .
Installing	Wait for 1 minute to install the ICAgent.
Installation failed	Uninstall the ICAgent by logging in to the server and then install it again.
Upgrading	Wait for 1 minute to upgrade the ICAgent.
Upgrade failed	Uninstall the ICAgent by logging in to the server and then install it again.
Offline	The AK/SK is incorrect or port 30200 or 30201 is disconnected. Solve the issue according to What Can I Do If an ICAgent Is Offline .
Faulty	The ICAgent is abnormal. Contact technical support.

- **Cause 2: AOM cannot monitor the current resource.**

Suggestion: Check whether your resources can be monitored by AOM. Specifically, AOM can monitor hosts, Kubernetes containers, and user processes, but cannot monitor system processes.

- **Cause 3: The resource is deleted or stopped.**

Suggestion:

- On the ECS page, check whether the host is restarted, stopped, or deleted.
- On the CCE page, check whether the service is stopped or deleted.
- If an application discovery rule is disabled or deleted, the application discovered based on the rule will also be disabled or deleted. Check whether the application discovery rule is disabled or deleted on AOM.

12.3.3 How Do I Set the Full-Screen Online Duration?

AOM provides an automatic logout mechanism to secure customer information. Specifically, after you access a page on the console but do not perform any operations within 1 hour, the console automatically logs you out.


If you set a full-screen view in the AOM console and later the system logs you out, the full-screen view will also exit. In this case, real-time monitoring cannot be performed. AOM allows you to customize full-screen online duration, meeting various requirements.

Precautions

- For security purposes, exit the full-screen view when it is not required.
- The full-screen online duration is irrelevant to operations. If the preset duration times out, the login page is automatically displayed.
- The full-screen online duration is subject to the last setting.
For example, if full-screen monitoring is implemented on multiple screens, the online duration is subject to the last setting.
For another example, if the online duration is set on both the **O&M** and **Dashboard** pages, the last setting prevails.
- The full-screen online duration takes precedence over the automatic logout mechanism of the cloud.
For example, if you log in to the console, set the full-screen online duration to 2 hours on AOM pages, and then open other pages, your setting on the AOM pages also takes effect on other pages. That is, the login page will be automatically displayed 2 hours later.
- If you leave all full-screen views, the default automatic logout mechanism is used.
For example, if you log in to the console, set the full-screen online duration to 2 hours on AOM pages, open other pages, and then leave all full-screen views of AOM, the default logout mechanism will be used. That is, if you do not perform any operations within 1 hour, the login page will be automatically displayed.

Setting the Full-Screen Online Duration on the Dashboard Page

Step 1 Log in to the AOM console. In the navigation pane, choose **Overview** > **Dashboard**.

Step 2 Click  in the upper right corner of the **Dashboard** page. In the dialog box that is displayed, set the full-screen online duration.

- **Custom:** The default online duration is 1 hour. You can enter 1–24 (unit: hour) in the text box.
For example, if you enter **2** in the text box, the login page is automatically displayed 2 hours later.
- **Always online:** The full-screen online duration is not restricted. That is, you can always implement full-screen monitoring and the login page will never be displayed.

Step 3 Click **OK** to enter the full-screen view.

----End

12.3.4 How Do I Obtain an AK/SK?

Each user can create a maximum of two Access Key ID/Secret Access Key (AK/SK) pairs. Once they are generated, they are permanently valid.

- AK: unique ID associated with the SK. It is used together with the SK to sign requests.
- SK: secret access key used in conjunction with an AK to sign requests cryptographically. It identifies a request sender and prevents the request from being modified.

Procedure

1. Log in to the management console, hover the mouse pointer over the username in the upper right corner, and select **My Credentials** from the drop-down list.
2. On the **My Credentials** page, click the **Access Keys** tab.
3. Click **Create Access Key** above the list and enter the verification code or password.
4. Click **OK** to download the generated AK/SK.

You can obtain the AK from the access key list and SK from the downloaded CSV file.

NOTE

- Keep the CSV file properly. You can only download the file right after the access key is created. If you cannot find the file, you can create an access key again.
- Open the CSV file in the lower left corner, or choose **Downloads** in the upper right corner of the browser and open the CSV file.
- Keep your access keys secure and change them periodically for security purposes.

12.3.5 How Can I Check Whether a Service Is Available?

Log in to the Application Operations Management (AOM) console, choose **Container Monitoring** in the navigation pane, and check the service status value at each time point in the workload monitoring view. If the value is **0**, the service is normal. Otherwise, the service is abnormal.

12.3.6 Why Is the Status of an Alarm Rule Displayed as "Insufficient"?

When you create an alarm rule for a resource, its data reported to AOM may be insufficient.

Possible causes:

1. The data reporting latency is too large. That is, the difference between the latest data reporting time of the line graph and the current time is greater than one threshold reporting period, which can be set to 1 minute or 5

minutes. If no data is obtained within such a period, a message indicating insufficient data is displayed.

2. If a metric is deleted or the host to which the metric belongs does not exist but the threshold rule still exists, a message indicating insufficient data is displayed.

12.3.7 Why the Status of a Workload that Runs Normally Is Displayed as "Abnormal" on the AOM Page?

A workload runs normally on Cloud Container Engine (CCE), but its status is **Abnormal** on the AOM page.

Possible causes:

1. The ICAgent version is too old.
ICAgent needs to be upgraded manually. If the ICAgent version is too old, the workload status may fail to be reported in time.
If the displayed workload status is incorrect, first check whether the ICAgent is in the latest version on the **Agent Management** page.
2. The node time is not synchronized with the actual time.
If the difference between the node time and the actual time is too large, the ICAgent fails to report metrics in time.
If the displayed workload status is incorrect, check whether the node time is different from the actual time or check the NTP offset on the AOM page.

12.3.8 How Do I Create the apm_admin_trust Agency?

Procedure

- Step 1** Log in to the IAM console.
- Step 2** In the navigation pane, choose **Agencies**.
- Step 3** On the page that is displayed, click **Create Agency** in the upper right corner. The **Create Agency** page is displayed.
- Step 4** Set parameters by referring to [Table 12-5](#).

Table 12-5 Parameters for creating an agency

Parameter	Description	Example
Agency Name	Set an agency name. NOTICE The agency name must be apm_admin_trust .	-
Agency Type	Select Cloud service .	Cloud service
Cloud Service	Select Application Operations Management (AOM) .	-

Parameter	Description	Example
Validity Period	Select Unlimited .	Unlimited
Description	(Optional) Provide details about the agency.	-

Step 5 Click **Next**. The **Authorize Agency** page is displayed.

Step 6 On the **Select Policy/Role** tab page, select **DMS UserAccess** and click **Next**.

DMS UserAccess: Common user permissions for DMS, excluding permissions for creating, modifying, deleting, scaling up instances and dumping.

Step 7 On the **Select Scope** tab page, set **Scope** to **Region-specific Projects** and select target projects under **Project [Region]**.

Step 8 Click **OK**.

----End

12.3.9 What Can I Do If an ICAgent Is Offline?

After an ICAgent is installed, its status is offline.

Problem Analysis

- **Cause:** The AK/SK are incorrect or ports 30200 and 30201 are disconnected.
- **Impact:** The ICAgent cannot work.

Solution

Step 1 Log in to the server where the ICAgent is installed as the **root** user.

Step 2 Run the following command to check whether the AK/SK configuration is correct:

```
cat /var/ICAgent/oss.icAgent.trace | grep proxyworkflow.go
```

- If no command output is displayed, the AK/SK configuration is incorrect. Go to [Step 3](#).
- If a command output is displayed, the AK/SK configuration is correct. Go to [Step 4](#).

Step 3 After configuring the AK/SK, reinstall the ICAgent. If the installation still fails, go to [Step 4](#).

Step 4 Check port connectivity.

1. Run the following command to obtain the access IP address:

```
cat /opt/oss/servicemgr/ICAgent/envs/ICProbeAgent.properties | grep ACCESS_IP
```

2. Run the following command to respectively check the connectivity of ports 30200 and 30201:

```
curl -k https://ACCESS_IP:30200
curl -k https://ACCESS_IP:30201
```

- If **404** is displayed, the port is connected. In this case, contact technical support.

- If the command output is not **404**, the port is not connected. Contact the network administrator to open the port and reinstall the ICAGENT. If the installation still fails, contact technical support.

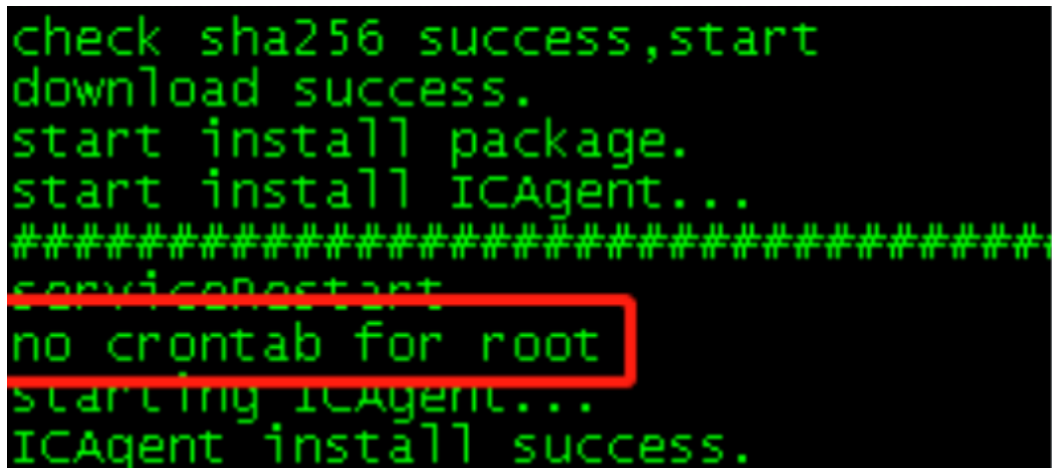
----End

12.3.10 Why Is "no crontab for root" Displayed During ICAGENT Installation?

Symptom

During ICAGENT installation, the system displays the message "no crontab for root".

Figure 12-2 Installing an ICAGENT



```
check sha256 success,start
download success.
start install package.
start install ICAGENT...
#####
serviceRestart
no crontab for root
starting ICAGENT...
ICAGENT install success.
```

Possible Causes

When you execute the script for installing an ICAGENT, crontab scheduled tasks will also be installed. The message "no crontab for root" indicates that there is no scheduled task of the **root** user.

Solution

No measure needs to be taken.

If the command output contains "ICAGENT install success", the ICAGENT is successfully installed and O&M data can be collected.

A Change History

Table A-1 Change history

Released On	Description
2024-06-30	<ul style="list-style-type: none"> • Added the following sections: <ul style="list-style-type: none"> - 3 Permissions Management - 6.1.4 Creating a Static Threshold Template - 6.1.5 Creating an Event Alarm Rule - 6.4 Alarm Action Rules - 6.5 Alarm Noise Reduction - 9.2.3 Setting Log Collection - 9.3 Quota Configuration - 9.4 Metric Configuration - 10 Auditing - 11 Upgrading to AOM 2.0 • Optimized the following sections: <ul style="list-style-type: none"> - 5.1 O&M - 6.1.3 Creating a Threshold Rule - 7.5 Metric Monitoring - 4.2 Configuring Application Discovery Rules - 2 Getting Started • Added sections 1 Service Overview, 4 Connecting Resources to AOM, and 12 FAQs. • Deleted sections "Creating a Notification Rule" and "Data Subscription."
2022-08-16	Optimized the description in 5.1 O&M .
2020-03-30	This issue is the first official release.