

Application Operations Management

User Guide

Issue 01
Date 2022-06-01



Copyright © Huawei Technologies Co., Ltd. 2022. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademarks and Permissions



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

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

Notice

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

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

Contents

1 Service Overview.....	1
1.1 What Is AOM?.....	1
1.2 Product Architecture.....	2
1.3 Functions.....	3
1.4 Metric Overview.....	4
1.4.1 Introduction.....	4
1.4.2 Network Metrics and Dimensions.....	5
1.4.3 Disk Metrics and Dimensions.....	6
1.4.4 File System Metrics and Dimensions.....	6
1.4.5 Host Metrics and Dimensions.....	7
1.4.6 Container Metrics and Dimensions.....	10
1.4.7 VM Metrics and Dimensions.....	12
1.4.8 Instance Metrics and Dimensions.....	13
1.4.9 Service Metrics and Dimensions.....	13
1.5 Restrictions.....	14
1.6 Glossary.....	18
1.7 Permissions Management.....	19
1.8 Billing.....	22
2 Getting Started.....	23
2.1 Process of Using AOM.....	23
2.2 Installing the ICAgent.....	25
2.3 Adding Alarm Rules and Viewing Alarms.....	25
3 User Guide.....	28
3.1 Overview.....	28
3.1.1 O&M.....	28
3.1.2 Dashboard.....	33
3.2 Alarm Management.....	39
3.2.1 Viewing Alarms.....	40
3.2.2 Viewing Events.....	40
3.2.3 Creating Alarm Rules.....	41
3.2.4 Creating Notification Rules.....	43
3.3 Resource Monitoring.....	44

3.3.1 Application Monitoring.....	45
3.3.2 Component Monitoring.....	46
3.3.3 Host Monitoring.....	47
3.3.4 Container Monitoring.....	48
3.3.5 Metric Monitoring.....	48
3.4 Log Management.....	51
3.4.1 Searching for Logs.....	51
3.4.2 Viewing Log Files.....	53
3.4.3 Configuring VM Log Collection Paths.....	55
3.5 Configuration Management.....	58
3.5.1 Agent Management.....	58
3.5.1.1 Installing the ICAgent.....	58
3.5.1.2 Upgrading the ICAgent.....	62
3.5.1.3 Uninstalling the ICAgent.....	63
3.5.2 Configuring Application Discovery.....	65
3.5.3 Log Configuration.....	70
3.5.3.1 Setting the Log Quota.....	70
3.5.3.2 Configuring Delimiters.....	71
4 FAQs.....	75
4.1 What Can I Do If an ICAgent Is Offline?.....	75
4.2 Obtaining an AK/SK.....	76
4.3 What Is the Relationship Between the Time Range and Statistical Cycle?.....	76
4.4 What Can I Do If Resources Are Not Running Properly?.....	77
4.5 How Can I Do If I Do Not Have the Permission to Access SMN?.....	79
4.6 How Do I Distinguish Alarms and Events?.....	79
4.7 Does AOM Display Logs in Real Time?.....	80
4.8 How Can I Check Whether a Service Is Available?.....	80
4.9 Why Is the Status of an Alarm Rule Displayed as "Insufficient"?.....	80
4.10 Why the Status of a Workload that Runs Normally Is Abnormal on the AOM Page?.....	81

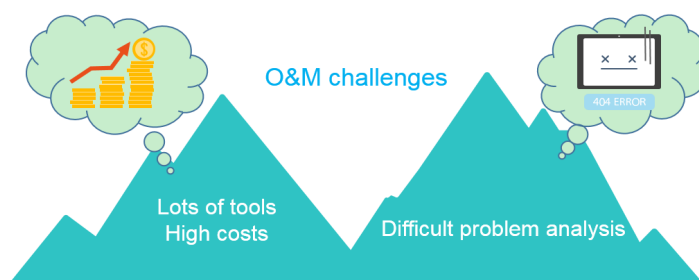
1 Service Overview

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 detect faults in a timely manner and monitor running status of applications, resources, and services in real time.

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 plug-in 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 through the open APIs or Exporter APIs provided by AOM.
- **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 a database. Cassandra stores sequential data, 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 > Container details > Process details. That is, applications, components, instances, containers, and processes are associated and their relationships are directly displayed 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, file systems, and alarms 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. By using the automatic discovery function, 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, digital 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 the same 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 List

Alarm management is to manage alarms and events.

You can create threshold rules for key resource metrics. When the metric data reaches the threshold, AOM generates alarms.

Log Management

AOM provides powerful log management capabilities. Log search enables you to quickly search for required logs from massive quantities of logs. By configuring delimiters, you can divide log content into multiple words and use these words to search for logs.

1.4 Metric Overview

1.4.1 Introduction

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

- System metrics: basic metrics provided by AOM, such as CPU usage and used CPU cores.
- Custom metrics: user-defined metrics. Custom metrics can be reported using the following methods:
 - Method 1: Use AOM APIs. For details, see Adding Monitoring Data and Querying Monitoring Data.
 - Method 2: Connect to Prometheus when creating containerized applications on the Cloud Container Engine (CCE) console. For details, see 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.

1.4.2 Network Metrics and Dimensions

Table 1-2 Network metrics

Metric	Description	Value Range	Unit
Downlink rate (recvBytesRate)	Inbound network traffic rate of a measured object	≥ 0	Byte Per Second (BPS)
Downlink rate (recvPackRate)	Number of data packets received by an NIC per second	≥ 0	Packet Per Second (PPS)
Downlink error rate (recvErrPackRate)	Number of error packets received by an NIC per second	≥ 0	PPS
Uplink rate (sendBytesRate)	Outbound network traffic rate of a measured object	≥ 0	BPS
Uplink error rate (sendErrPackRate)	Number of error packets sent by an NIC per second	≥ 0	PPS
Uplink rate (sendPackRate)	Number of data packets sent by an NIC per second	≥ 0	PPS
Total rate (totalBytesRate)	Total inbound and outbound network traffic rate of a measured object	≥ 0	BPS

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

1.4.3 Disk Metrics and Dimensions

Table 1-4 Disk metrics

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

1.4.4 File System Metrics and Dimensions

Table 1-6 File system metrics

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

Metric	Description	Value Range	Unit
Disk read/write status (diskRWStatus)	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 (diskUsedRate)	Percentage of the used disk space to the total disk space	≥ 0	%

Table 1-7 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

1.4.5 Host Metrics and Dimensions

Table 1-8 Host metrics

Metric	Description	Value Range	Unit
Total CPU cores (cpuCoreLimit)	Total number of CPU cores that have been applied for a measured object	≥ 1	Cores
Used CPU cores (cpuCoreUsed)	Number of CPU cores used by a measured object	≥ 0	Cores

Metric	Description	Value Range	Unit
CPU usage (cpuUsage)	CPU usage of a measured object	0%–100%	%
Available physical memory (freeMem)	Available physical memory of a measured object	≥ 0	MB
Available virtual memory (freeVirMem)	Available virtual memory of a measured object	≥ 0	MB
Total GPU memory (gpuMemCapacity)	Total GPU memory of a measured object	> 0	MB
GPU memory usage (gpuMemUsage)	Percentage of the used GPU memory to the total GPU memory	0%–100%	%
Used GPU memory (gpuMemUsed)	GPU memory used by a measured object	≥ 0	MB
GPU usage (gpuUtil)	GPU usage of a measured object	0%–100%	%
Physical memory usage (memUsedRate)	Percentage of the used physical memory to the total physical memory	0%–100%	%
Host status (nodeStatus)	Host status	<ul style="list-style-type: none"> • 0: Normal • 1: Abnormal 	N/A
NTP offset (ntpOffset)	Offset between the local time of the host and the NTP server time. When the NTP offset is closer to 0, the local time of the host is closer to the time of the NTP server.	N/A	ms
NTP server status (ntpServerStatus)	Whether the host is connected to the NTP server	0 or 1 <ul style="list-style-type: none"> • 0: Connected • 1: Unconnected 	N/A

Metric	Description	Value Range	Unit
NTP sync status (ntpStatus)	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: Asynchronous 	N/A
Processes (processNum)	Number of processes on a measured object	≥ 0	N/A
GPU temperature (temperature)	GPU temperature of a measured object	-	°C
Total physical memory (totalMem)	Total physical memory that has been applied for a measured object	≥ 0	MB
Total virtual memory (totalVirMem)	Total virtual memory that has been applied for a measured object	≥ 0	MB
Virtual memory usage (virMemUsedRate)	Percentage of the used virtual memory to the total virtual memory	0%–100%	%
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

 **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)

Currently, the virtual memory of a newly created 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.

Table 1-9 Dimensions of host metrics

Dimension	Description
clusterId	Cluster ID
clusterName	Cluster name
gpuName	GPU name
gpuID	GPU ID
hostID	Host ID
nameSpace	Cluster namespace
nodeIP	Host IP address
nodeName	Host name

1.4.6 Container Metrics and Dimensions

Table 1-10 Container metrics

Metric	Description	Value Range	Unit
Total CPU cores (cpuCoreLimit)	Total number of CPU cores that have been applied for a measured object	≥ 1	Cores
Used CPU cores (cpuCoreUsed)	Number of CPU cores used by a measured object	≥ 0	Cores
CPU usage (cpuUsage)	CPU usage of a measured object Percentage of the used CPU cores to the total CPU cores	0%–100%	%
Disk read rate (diskReadRate)	Volume of data read from a disk per second	≥ 0	KB/s
Disk write rate (diskWriteRate)	Volume of data written into a disk per second	≥ 0	KB/s
Available file system (filesystemAvailable)	Available file system capacity of a measured object Only containers using the device mapper in the Kubernetes cluster of version 1.11 or later are supported.	≥ 0	MB
Total file system (filesystemCapacity)	Total file system capacity of a measured object Only containers using the device mapper in the Kubernetes cluster of version 1.11 or later are supported.	≥ 0	MB

Metric	Description	Value Range	Unit
File system usage (filesystemUsage)	File system usage of a measured object Percentage of the used file system to the total file system Only containers using the device mapper in the Kubernetes cluster of version 1.11 or later are supported.	0%–100%	%
Total physical memory (memCapacity)	Total physical memory that has been applied for a measured object	≥ 0	MB
Physical memory usage (memUsage)	Percentage of the used physical memory to the total physical memory	0%–100%	%
Used physical memory (memUsed)	Used physical memory of a measured object	≥ 0	MB
Downlink rate (recvBytesRate)	Inbound network traffic rate of a measured object	≥ 0	Byte Per Second (BPS)
Downlink rate (recvPackRate)	Number of data packets received by an NIC per second	≥ 0	Packet Per Second (PPS)
Downlink error rate (recvErrPackRate)	Number of error packets received by an NIC per second	≥ 0	PPS
Error packets (rxPackErrors)	Number of error packets received by a measured object	≥ 0	Packets
Uplink rate (sendBytesRate)	Outbound network traffic rate of a measured object	≥ 0	BPS
Uplink error rate (sendErrPackRate)	Number of error packets sent by an NIC per second	≥ 0	PPS
Uplink rate (sendPackRate)	Number of data packets sent by an NIC per second	≥ 0	PPS
Status (status)	Docker container status	0 or 1 <ul style="list-style-type: none"> • 0: Normal • 1: Abnormal 	N/A

Table 1-11 Dimensions of container metrics

Dimension	Description
appID	Service ID

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

1.4.7 VM Metrics and Dimensions

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

Table 1-12 Process metrics

Metric	Description	Value Range	Unit
Total CPU cores (cpuCoreLimit)	Total number of CPU cores that have been applied for a measured object	≥ 1	Cores
Used CPU cores (cpuCoreUsed)	Number of CPU cores used by a measured object	≥ 0	Cores
CPU usage (cpuUsage)	CPU usage of a measured object Percentage of the used CPU cores to the total CPU cores	0%–100%	%
Handles (handleCount)	Number of handles used by a measured object	≥ 0	N/A
Total physical memory (memCapacity)	Total physical memory that has been applied for a measured object	≥ 0	MB
Physical memory usage (memUsage)	Percentage of the used physical memory to the total physical memory	0%–100%	%
Used physical memory (memUsed)	Used physical memory of a measured object	≥ 0	MB

Metric	Description	Value Range	Unit
Status (status)	Process status	0 or 1 <ul style="list-style-type: none"> 0: Normal 1: Abnormal 	N/A
Threads (threadsCount)	Number of threads used by a measured object	≥ 0	N/A
Total virtual memory (virMemCapacity)	Total virtual memory that has been applied for a measured object	≥ 0	MB

Table 1-13 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

1.4.8 Instance Metrics and Dimensions

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

1.4.9 Service Metrics and Dimensions

Service metrics consist of instance metrics. The dimensions of service metrics are the same as those of instance metrics. For details, see [Instance Metrics and Dimensions](#).

1.5 Restrictions

OS Usage Restrictions

AOM supports multiple operating systems (OSs). When creating a host, ensure that its OS meets the requirements in [Table 1-14](#). Otherwise, the host cannot be monitored by AOM.

Table 1-14 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			
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	

 **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 1-15](#).

Table 1-15 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 can be added to a line graph, and resources can be selected across clusters. • Only one resource can be added to a digital graph. • 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.
	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.
	Data points returned for a single query	A maximum of 1440 data points can be returned each time.
	Custom metric	No restrictions.
	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.


Category	Object	Usage Restrictions
	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 exceeds 10 KB, the ICAgent does not collect it. That is, 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.
		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.

Category	Object	Usage Restrictions
	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. It is recommended that the log generation speed of a single node be lower than 5 MB/s.
	Log loss	When a single log line exceeds 10,240 bytes, the line will be discarded.
	Log repetition	When the ICAgent is restarted, identical data may be collected around the restart time.
Alarm center	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, as shown in [Figure 1-3](#).

Figure 1-3 Changing the interpolation mode



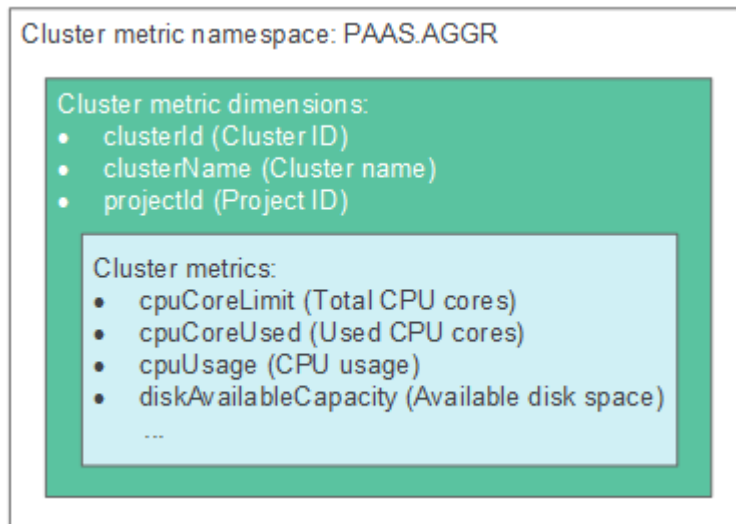
1.6 Glossary

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 1-4** describes the relationships among namespaces, dimensions, and cluster metrics.

Figure 1-4 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, ECS) that you created. A host can be connected to AOM for monitoring as long as its OS is supported by AOM and 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 search and analysis of massive quantities of logs.

Alarms

Alarms are reported when AOM or an external service such as ServiceStage, Cloud Container Engine (CCE), or Application Performance Management (APM) 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, CCE, or APM encounters some changes. Such changes are not necessarily cause service exceptions. Events do not need to be handled.

Threshold Rules

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

1.7 Permissions Management

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 must not 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 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.

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.

Policies are 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 1-16 lists all the system permissions supported by AOM.

Table 1-16 System permissions supported by AOM

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

Table 1-17 lists the common operations supported by each system-defined policy of AOM. Please choose proper system-defined policies according to this table.

Table 1-17 Common operations supported by each system-defined policy of AOM

Operation	AOM FullAccess	AOM ReadOnlyAccess
Creating a threshold rule	√	x
Modifying a threshold rule	√	x
Deleting a threshold rule	√	x
Creating a threshold template	√	x
Modifying a threshold template	√	x
Deleting a threshold template	√	x
Creating a dashboard	√	x
Modifying a dashboard	√	x
Deleting a dashboard	√	x
Creating a notification rule	√	x
Modifying a notification rule	√	x
Deleting a notification rule	√	x
Creating an application discovery rule	√	x
Modifying an application discovery rule	√	x
Deleting an application discovery rule	√	x
Subscribing to threshold alarms	√	x
Exporting a monitoring report	√	√
Configuring a VM log collection path	√	x
Adding a log bucket	√	x
Modifying a log bucket	√	x
Deleting a log bucket	√	x
Adding an extraction rule	√	x

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

1.8 Billing

Billing

 NOTE

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

2 Getting Started

2.1 Process of Using AOM

AOM is a one-stop, multi-dimensional O&M management platform for cloud applications. It monitors applications and related cloud resources in real time, analyzes application health status, and provides flexible alarm reporting and data visualization functions. It helps you detect faults in a timely manner and monitor running status of applications, services, and other resources in real time. This section describes how to get started with AOM. The following figure shows the process.

Figure 2-1 Process of using AOM



1. Creating a cloud host
Each host corresponds to a VM on the cloud, for example, an Elastic Cloud Server (ECS). A host can be directly created on the ECS console, or indirectly created on the Cloud Container Engine (CCE) console.
2. Installing the ICAgent
ICAgent is the data collector of AOM. It collects metrics, logs, and application performance data in real time. For hosts created on the ECS console, you need to manually install the ICAgent. For hosts created on the CCE console, the ICAgent is automatically installed.
3. Configuring an alarm rule
You can set threshold conditions for metrics by using alarm rules. If a metric value meets a threshold condition, AOM generates a threshold alarm. If no metric data is reported, AOM will report an insufficient data event. In this way, you can identify and handle exceptions at the earliest time.
4. Viewing alarms

AOM provides the dashboard and alarm list for you to perform routine O&M.

2.2 Installing the ICAgent

This section describes how to install an ICAgent on an ECS.

Prerequisites

- An ECS has been created.
- An EIP has been bound to the ECS.
- An AK/SK have been obtained. For details, see [Obtaining an AK/SK](#).
- The browser time is the same as the ECS time.

Procedure

- Step 1** Log in to the AOM console and choose **Configuration Management > Agent Management** in the navigation pane.
- Step 2** Click **Install ICAgent**.
- Step 3** Generate and copy the ICAgent installation command.
- Step 4** Use a remote login tool to log in as the **root** user to the server where the ICAgent is to be installed, run the command copied in the previous step, and enter the AK/SK as prompted to install the ICAgent.

NOTE

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

----End

2.3 Adding Alarm Rules and Viewing Alarms

You can set threshold conditions for metrics by using alarm rules. If metric values meet threshold conditions, AOM will generate threshold alarms. If no metric data is reported, AOM will report insufficient data events. In this way, you can identify and handle exceptions at the earliest time.

For example, during routine O&M, a host may break down or restart due to an excessively-high CPU usage. To avoid the problem, set an alarm rule. For example, when the CPU usage of a host exceeds 85%, an alarm is reported, so that you can quickly obtain the resource running status and take measures to prevent service loss. This section describes how to add alarm rules and view alarms.

Procedure

- Step 1** In the navigation pane, choose **Alarm Center > Alarm Rule** and click **Add Threshold**.

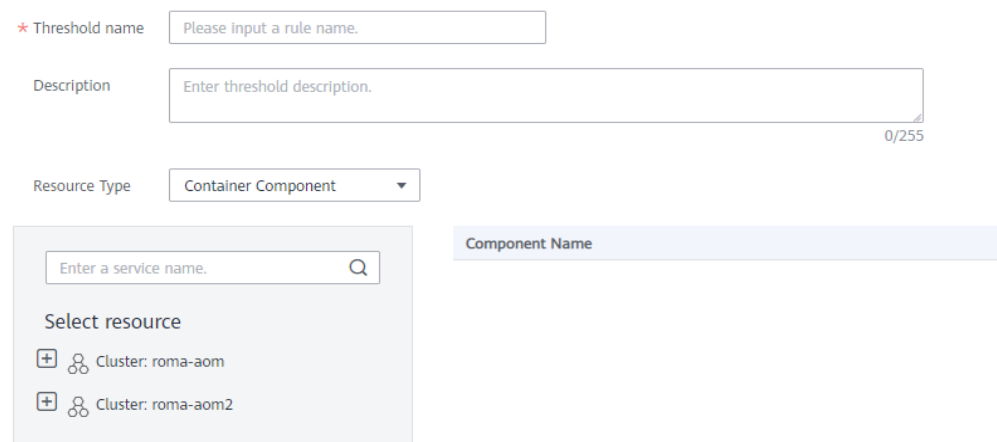
Step 2 Customize alarm rules.

1. Select resources: Enter a rule name, select a resource type, select the resources to be monitored from the resource tree, and click **Next**.

NOTE

- You can select a maximum of 100 resources from the resource tree.
- When multiple resources are selected, multiple alarm rules will be created after the creation is complete. Each resource is monitored by an alarm rule. A rule name consists of the alarm rule name you enter in the **Threshold name** text box, and a sequence number ranging from 0 to 99. The earlier a resource is selected, the smaller its sequence number.

Figure 2-2 Selecting resources



2. Customize a threshold: Select the metric to be monitored, and set parameters such as **Threshold Condition**, **Consecutive Period (s)**, **Alarm Severity**, and **Statistic Method**.


NOTE

- **Threshold Condition:** Trigger condition of a threshold alarm. A threshold condition consists of two parts: determination condition (\geq , \leq , $>$, or $<$) and threshold value. For example, if **Threshold Condition** is set to > 85 and an actual metric value exceeds 85, a threshold alarm will be generated.
- **Consecutive Period (s):** If a metric value meets the threshold condition for a specified number of consecutive periods, a threshold alarm will be generated.
- **Statistic Method:** Method used to measure metric values.
- **Statistical Cycle:** Interval at which metric data is collected.

- Step 3** Click **Submit**. As shown in the following figure, multiple alarm rules are created. Each resource is monitored by an independent rule.

For example, when a monitored component uses more than 3 CPU cores, a threshold alarm is generated on the alarm page. You can choose **Alarm Center > Alarm Rule** in the navigation pane to view the alarm rule.

Figure 2-3 Alarm rules




Step 4 In the navigation pane, choose **Alarm Center > Alarm List**.

Step 5 View alarms on the **Alarm List** page.

- Set a time range to view alarms. There are two methods to set a time range:
 Method 1: Use a predefined time label, such as **Last 1 hour**, **Last 6 hours**, or **Last 1 day** in the upper right corner of the page. You can select a time range as required.
 Method 2: Specify the start time and end time in the upper right corner of the page to customize a time range. You can specify up to 30 days.
- Set filter criteria and click **Search** to view the alarms generated in the specified time range.

Step 6 Perform the operations listed in [Table 2-1](#) as required.

Table 2-1 Operations

Operation	Method	Description
Viewing alarm statistics	View alarm statistics that meet filter criteria within a specific time range through 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 an alarm after the corresponding problem is resolved. You can view the cleared alarms on the History tab page.
Viewing alarm details	Click an alarm name to view alarm details.	-

----End

3 User Guide

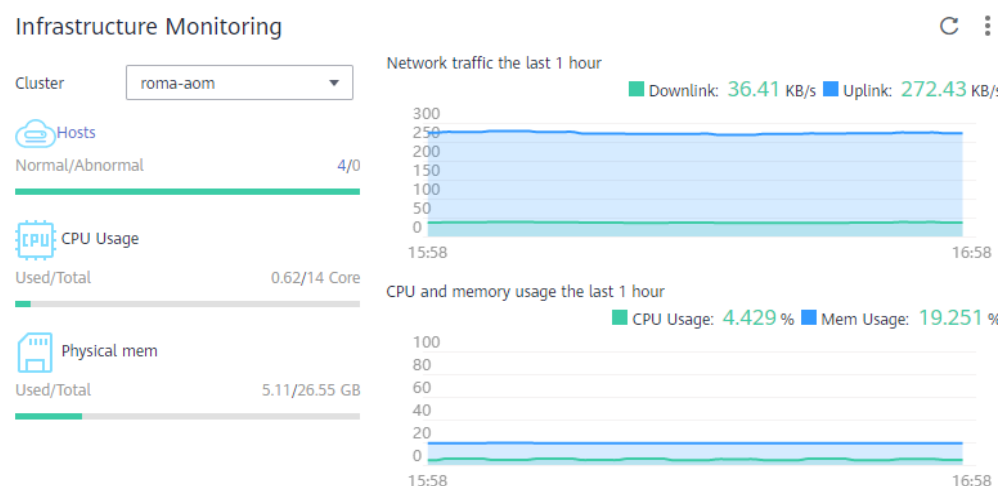
3.1 Overview

3.1.1 O&M

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

Infrastructure Monitoring

Figure 3-1 Infrastructure monitoring

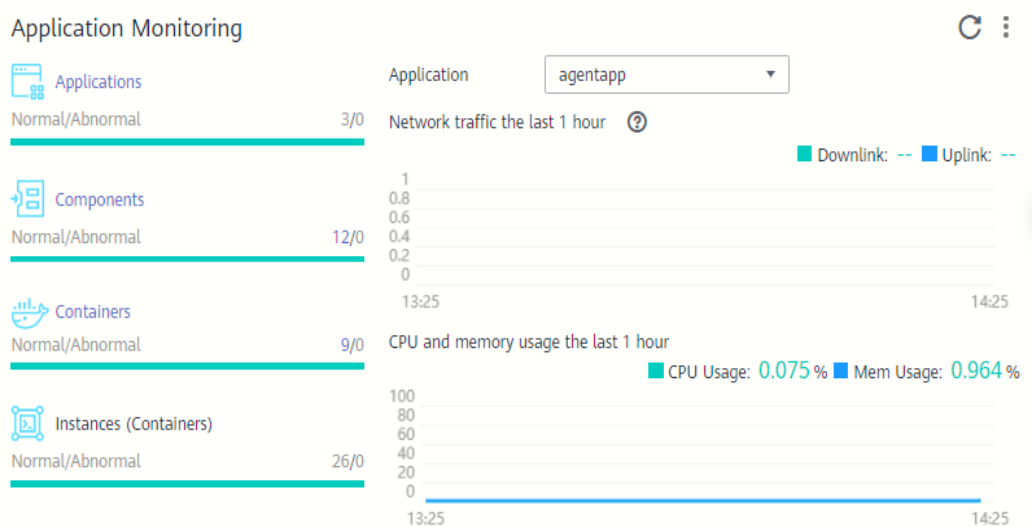


This card mainly displays infrastructure metrics. You can select one cluster to view its information. When you select the **roma-aom2** cluster, the following information is displayed:

- Host running status, CPU usage, and physical memory usage.
- Trend graph of network traffic data in the last hour. The values of each point in the graph respectively indicate the total downlink and uplink traffic of all clusters in one minute. The values above the graph respectively indicate the total downlink and uplink traffic of the cluster at the latest time point.
- Trend graph of CPU and memory usage in the last hour. The values of each point in the graph respectively indicate the average CPU and memory usage of the cluster in one minute. The values above the graph respectively indicate the average CPU and memory usage of the cluster at the latest time point.

Application Monitoring

Figure 3-2 Application monitoring

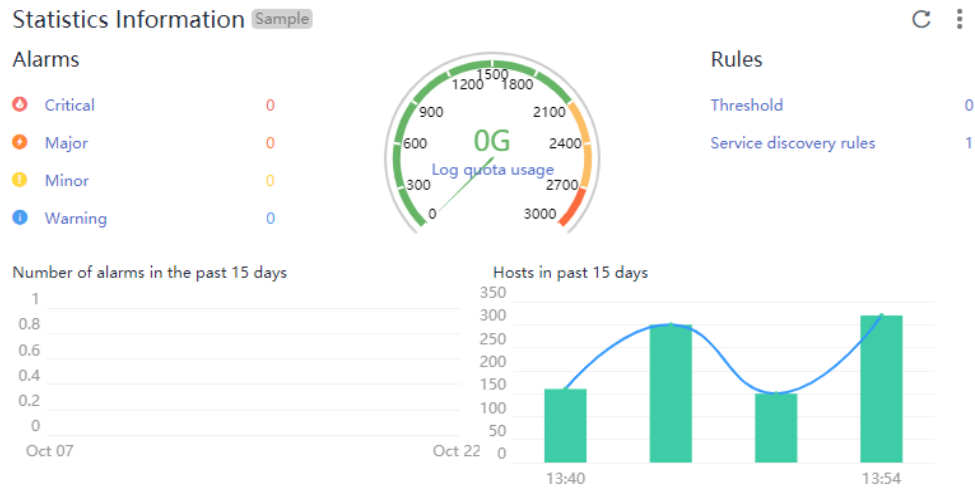


This card mainly displays application metrics:

1. Running status of applications, components, containers, and instances.
2. The following information is displayed when you select an application:
 - Trend graph of network traffic data in the last hour. 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 hour. 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.

Information Statistics


Figure 3-3 Information statistics



This card mainly displays alarms, alarm rules, and trends of alarms and hosts.

Host Monitoring (Disk)

Figure 3-4 Host monitoring (disk)




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 average 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.

Cluster Monitoring (CPU and Memory)

Figure 3-5 Cluster monitoring (CPU and memory)

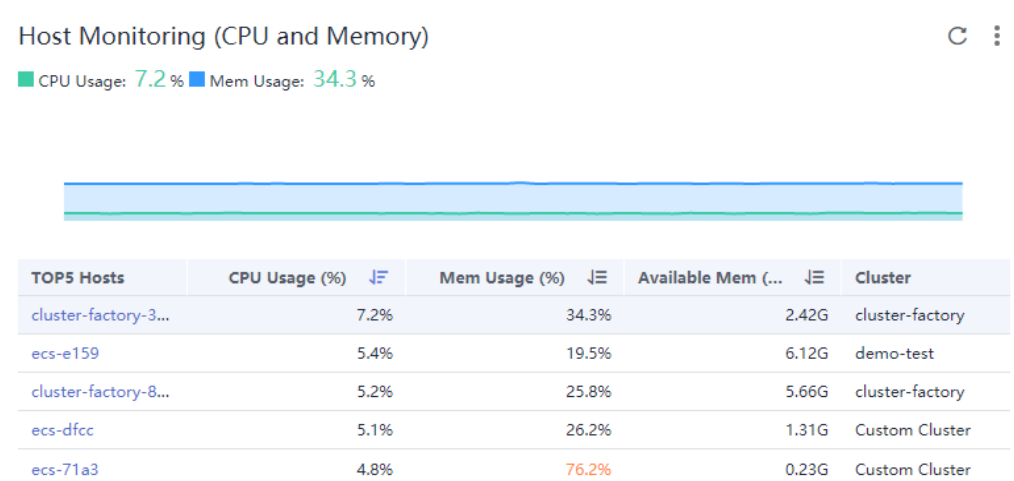


This card mainly displays:

- 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 average 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.

Host Monitoring (CPU and Memory)

Figure 3-6 Host monitoring (CPU and memory)



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 average 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.


Container Instance Monitoring (CPU and Memory)

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 average 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.
- option, which can be selected as required.

Host Monitoring (Network)

Figure 3-7 Host monitoring (network)




This card mainly displays:

- The top 5 hosts with high uplink/downlink network traffic in the last minute.
- Trend graph of the uplink/downlink network traffic of the selected host in the last hour. The values of each point in the graph respectively indicate the average uplink/downlink network traffic of the selected host in one minute.
- Uplink/downlink network traffic of the selected host at the latest time point, which is displayed above the trend graph.

Cluster Monitoring (Disk)

Figure 3-8 Cluster monitoring (disk)





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 average 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.

More Operations

You can also perform the operations described in [Table 3-1](#).

Table 3-1 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.
Enlarging a 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.

3.1.2 Dashboard

With a dashboard, different graphs can be displayed on the same screen. Different graphs, such as line graphs and digit graphs can be used to display resource data, which lets you view monitoring data comprehensively.

For example, you can add key metrics of important resources to the dashboard for real-time monitoring. You can also compare the same metric of different resources on one screen. 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. The following figure shows the CPU cores of different components.

Figure 3-9 Line graph



- **Digit graph:** displays the latest value of a metric in real time. The following figure shows the average CPU usage of a component.

Figure 3-10 Digit graph



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

The statuses of thresholds, hosts, and components can be displayed. The statuses of one or more alarm rules, hosts, or components can be added in one graph for monitoring.

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

Figure 3-11 Threshold-crossing status graph

Threshold Status		Operation ▾
Threshold Name	Threshold Status	
cputest0627	Exceeded	
aom-test	Insufficient	

NOTE

Before adding a threshold-crossing status graph, [create an alarm rule](#). Otherwise, the threshold-crossing status graph cannot be added.


- **Host status graph:** monitors the health status of hosts in real time.

Figure 3-12 Host status graph

Host status			Operation ▾
Host Name	Alias	Status	
roma-aom-07972-hw7pj	--	Normal	

- **Component status graph:** monitors the health status of components in real time.

Figure 3-13 Component status graph




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 the top N hosts for display. If the number of resources is less than N, actual resources are displayed. The following figure lists the top 5 instances with the highest CPU usage.

Figure 3-14 Top N resource graph



NOTE

- By default, the top 5 resources are displayed. To view the top 15 resources, click **Show Top15**, 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.
- 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 100 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 alarm 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**.

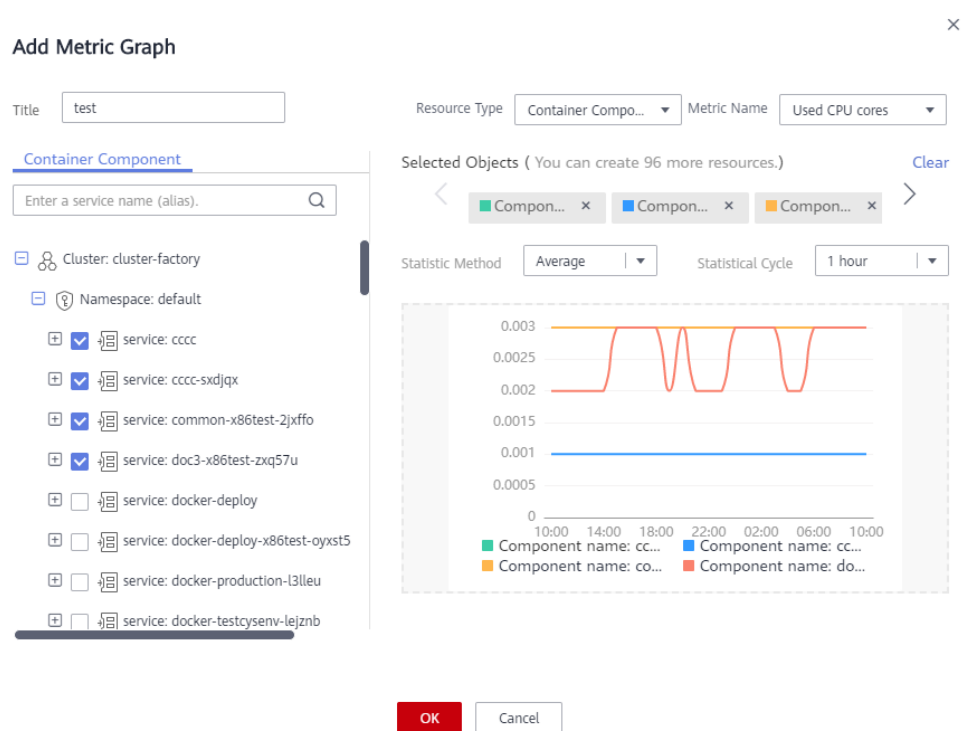
Step 3 Click **Add Metric Graph** in the upper right corner of the **Dashboard** page. Then, select **Metric Data**, **Health Status**, or **Top N Resources**.

Step 4 Add a metric graph to the dashboard.

- Under **Metric Data**, line and digit graphs can be selected.
- Under **Health Status**, threshold-crossing status graphs, host status graphs, and component status graphs can be selected.


Select a graph that is appropriate for 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 **Statistic Method** and **Statistical Cycle**, and click **OK**.



Step 5 Click **Save** in the upper right corner of the **Dashboard** page.

NOTE

The **Auto Refresh** () option in the upper right corner of the **Dashboard** page is used to automatically refresh all graphs in the dashboard.

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

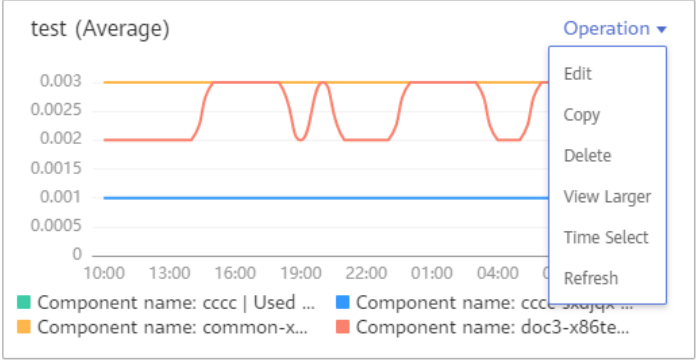

----End

More Operations

After creating a dashboard, you can also perform the operations described in [Table 3-2](#).

Table 3-2 Related operations

Operation 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	

Operation Object	Operation	Description
	Export monitoring report	Click Export Monitoring Report to export line graphs in the dashboard as CSV files to a local PC.
Graph	Add	Click Add Metric Graph to add a line graph, digit graph, threshold-crossing status graph, host status graph, component status graph, or top N resource graph to the dashboard.
	Edit	Choose Edit , Copy , Delete , or View Larger (support only for line graphs) from the Operation drop-down list. 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.
	Copy	
	Delete	
	Zoom in	
	Time select	
	Refresh	<p>Figure 3-15 Operations on a graph</p>  <p>NOTE 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>
Resize	Hover over the lower right corner of a graph. When the cursor changes to  , hold down your left mouse button to resize the graph.	
Reposition	Hover over the blank area in the upper or lower part of a graph, and drag and drop it to the desired position.	

3.2 Alarm Management


3.2.1 Viewing Alarms

Alarms are the information which is reported when AOM or an external service is abnormal or may cause exceptions. You need to take measures to resolve faults. Otherwise, service exceptions may occur.

Viewing Alarms


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

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

- Set a time range to view alarms. There are two methods to set a time range:
 Method 1: Use a predefined time label, such as **Last 1 hour**, **Last 6 hours**, or **Last 1 day** in the upper right corner of the page. You can select a time range as required.
 Method 2: Specify the start time and end time in the upper right corner of the page to customize a time range. You can specify up to 30 days.
- Set filter criteria and click  to view the alarms generated in the specified time range.

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

Table 3-3 Operations

Operation	Method	Description
Viewing alarm statistics	View alarm statistics that meet filter criteria within a specific time range through a bar graph.	-
Clearing alarms	In the alarm list, click  in the Options column of the target alarm.	<ul style="list-style-type: none"> You can clear an alarm after the corresponding problem is resolved. You can view the cleared alarms on the History tab page.
Viewing alarm details	Click an alarm name to view alarm details.	-

----End


3.2.2 Viewing Events

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

Viewing Events

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

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

1. Set a time range to view events. There are two methods to set a time range:
 Method 1: Use a predefined time label, such as **Last 1 hour**, **Last 6 hours**, or **Last 1 day** in the upper right corner of the page. You can select a time range as required.
 Method 2: Specify the start time and end time in the upper right corner of the page to customize a time range. You can specify up to 30 days.
2. Set filter criteria and click  to view the events generated in the specified time range.
 You can click **x** to clear filter criteria.

Step 3 Perform the operations listed in [Table 3-4](#) as required.

Table 3-4 Operations

Operation	Method	Description
Viewing event statistics	View event statistics that meet filter criteria within a specific time range through a bar graph.	-

----End

3.2.3 Creating Alarm Rules

You can set threshold conditions for metrics by using alarm rules. If metric values meet threshold conditions, AOM will generate threshold alarms. If no metric data is reported, AOM will report insufficient data events. In this way, you can identify and handle exceptions at the earliest time.

For example, during routine O&M, a host may break down or restart due to an excessively-high CPU usage. To avoid the problem, set an alarm rule. For example, when the CPU usage of a host exceeds 85%, an alarm is reported, so that you can quickly obtain the resource running status and take measures to prevent service loss.

Precautions

You can create a maximum of 1000 alarm rules. If the number of alarm rules reaches the upper limit, delete unnecessary rules and create new ones.

Customizing Alarm Rules

Step 1 In the navigation pane, choose **Alarm Center > Alarm Rule** and click **Add Threshold**.

Step 2 Customize alarm rules.

1. Select a resource. Specifically, enter a threshold rule name, select a resource type, select the resource to be monitored from the resource tree, and click **Next**.

NOTE

- You can select a maximum of 100 resources from the resource tree.
- When multiple resources are selected, multiple alarm rules will be created after the creation is complete. Each resource is monitored by an alarm rule. A rule name consists of the alarm rule name you enter in the **Threshold name** text box, and a sequence number ranging from 0 to 99. The earlier a resource is selected, the smaller its sequence number.

Figure 3-16 Selecting resources

The screenshot displays the configuration interface for creating an alarm rule. It includes the following elements:

- Threshold name:** A text input field with the placeholder text "Please input a rule name."
- Description:** A larger text input field with the placeholder text "Enter threshold description." and a character count "0/255" at the bottom right.
- Resource Type:** A dropdown menu currently set to "Container Component".
- Resource Selection:** A dropdown menu showing a search bar "Enter a service name." and a list of resources under the heading "Select resource". The visible resources are "Cluster: roma-aom" and "Cluster: roma-aom2", each with a plus icon and a cluster icon.
- Component Name:** A light blue header bar for the resource selection dropdown.

2. Customize a threshold. Specifically, select the metric to be monitored, and set parameters such as **Threshold Condition**, **Consecutive Period (s)**, **Alarm Severity**, and **Statistic Method**.

NOTE

- **Threshold Condition:** Trigger condition of a threshold alarm. A threshold condition consists of two parts: determination condition (\geq , \leq , $>$, or $<$) and threshold value. For example, if **Threshold Condition** is set to > 85 and an actual metric value exceeds 85, a threshold alarm will be generated.
- **Consecutive Period (s):** If a metric value meets the threshold condition for a specified number of consecutive periods, a threshold alarm will be generated.
- **Alarm Severity: Critical, Major, Minor, or Warning.**
- **Send Notification:** Alarm notifications sent by SMN will be billed according to the standard pricing of SMN. If you select **Yes**, select a topic and trigger condition.
 - **Topic:** If the topic you want to select does not exist, create one on the SMN console.
 - **Trigger Condition:** The condition can be **Threshold crossing, Normal, or Insufficient data.**
- **Statistic Method:** Method used to measure metric values.
- **Statistical Cycle:** Interval at which metric data is collected.

Step 3 Click **Submit**. As shown in the following figure, multiple alarm rules are created. Each resource is monitored by an independent rule.

For example, when a monitored component uses more than 3 CPU cores, a threshold alarm is generated on the alarm page. You can choose **Alarm Center > Alarm Rule** in the navigation pane to view the alarm rule.

Figure 3-17 Alarm rules

Name	Status	Metric Name	Operation
test1	OK	Used CPU cores	Modify Delete
test2	OK	Used CPU cores	Modify Delete
test3	OK	Used CPU cores	Modify Delete
test0	Insufficient	Used CPU cores	Modify Delete

----End

3.2.4 Creating Notification Rules

Application Operations Management (AOM) supports alarm notification. You can use this function by creating notification rules. When alarms are reported due to an exception in AOM or an external service, alarm information can be sent to specified personnel by email or Short Message Service (SMS) message. In this way, these personnel can rectify faults in time to avoid service loss.

If no notification rules exist, no alarm notifications will be sent. In this case, you can only view alarms on the **Alarm List** page in the AOM console.

Procedure

After notification rules are created, SMS messages or emails are sent when the notification rules are met.

For example, to ensure that O&M personnel can receive notifications by email when the ICAgent on HostA, HostB, HostC, or HostD in Cluster1 is abnormal or fails to be installed, upgraded, or uninstalled, perform the following operations:

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

Step 2 Click **Create SMN Topic** and set a notification policy on the Simple Message Notification (SMN) console when AOM is interconnected with SMN. If you have configured a notification policy, skip this step.

1. Create a topic according to "Creating a Topic" in the *SMN User Guide*.
For example, create a topic named **Topic1**.
2. Set a topic policy according to "Creating a Topic" in the *SMN User Guide*.
Select **apm** for **Services that can publish messages to this topic**. Otherwise, notifications will fail to be sent.
3. Add a subscriber, that is, the email or SMS message recipient, for the topic according to "Adding a Subscription" in the *SMN User Guide*. In this way, SMN can notify subscribers of alarm information in real time.
For example, enter O&M personnel's email addresses.

Step 3 Create a notification rule. Specifically, enter the rule name, select the notification condition, select the topic created in **Step 2**, select the time zone and language as required, enter the rule description, and click **Confirm**.


After a notification rule is created, the O&M personnel will receive an email or SMS notification when this rule is met.

----End

More Operations

After creating a notification rule, you can also perform the operations described in [Table 3-5](#).

Table 3-5 Related operations

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

3.3 Resource Monitoring

3.3.1 Application Monitoring

An application is a group of identical or similar components divided based on business requirements. Applications are classified into system applications and custom applications. System applications are discovered based on built-in discovery rules, and custom applications are discovered based on custom rules.

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

Procedure

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

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

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

Step 3 On the application details page, click the **View Monitor Graphs** tab and monitor application metrics.

You can also perform the following operations:


- **Adding an application**

For the same or similar components that are discovered by default discovery rules or that are not installed with 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** to add a custom application discovery rule. For details, see [Configuring Application Discovery](#). After the rule is added, you can monitor the application. AOM can display O&M information by component. For details, see [Component Monitoring](#).

- **Creating a view template**



AOM provides default view templates, such as **Application Template**. You can modify the default view templates, and can also click the plus sign (+) in



to create custom templates.

- **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 according to [Dashboard](#).

- **Adding a view template to a dashboard**

On the application details page, choose **More > Add To Dashboard** in the upper right corner to add the view template to the dashboard for monitoring.

----End


3.3.2 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 information such as type, CPU usage, and memory usage of each component. 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. You can implement multi-dimensional monitoring.

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

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

- Click  in the upper right corner and select **Hide system component**.

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.

- **Adding a tag**


Tags are used to identify components. You can distinguish system components from non-system ones by using 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 and select or deselect **Hide system component**. You can also customize tags to facilitate component management.

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

that contains the target component, click , enter a tag, and click  and **OK**.

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 them.

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

Step 4 Click the component name to go to the **Component Details** page.

- Click **View Log** next to the component name to go to the log search page and view the logs of the component. Logs of ServiceStage components cannot be viewed on the **Component Details** page.
- On the **Instance List** tab page, view the instance details.

 **NOTE**

Click an instance name to monitor the resource usage and health status of service processes or instances.

- 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.

----End

3.3.3 Host Monitoring

Hosts include Elastic Cloud Servers (ECSs). AOM monitors the hosts created during Cloud Container Engine (CCE) or ServiceStage cluster creation and the hosts that are directly created. Ensure that the directly created hosts meet operating system (OS) and version requirements and install the ICAgent on these hosts according to [Installing the ICAgent](#). Otherwise, these hosts cannot be monitored by AOM. In addition, the hosts support both IPv4 and IPv6 addresses.

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

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 when the host generates a threshold alarm. For more information, see [What Can I Do If Resources Are Not Running Properly?](#)

Procedure

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

Click  in the upper right corner and select **Hide master host**.

Step 2 Perform the following operations as required:



- **Adding an alias**

If a host name is too complex to identify, you can add an alias that is easy to identify a host as required.

Click **Add alias** in the **Operations** column of the target host.

- **Adding a tag**

Tags are identifiers of hosts. You can manage hosts using tags. After a tag is added, you can quickly identify and select a host.

In the host list, choose **More > Add tags** in the **Operation** column, enter a tag, and click  and **OK**. 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 them.

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

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

Step 5 Monitor common system devices such as GPUs and NICs 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 GPUs of the host. Click a GPU to view its metrics on the **View Monitor Graphs** page.
- Click the **NIC** tab to view the basic information about the NICs 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 disks 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.

----End

3.3.4 Container Monitoring

The difference between container and component monitoring is that their monitored objects are different. For component monitoring, workloads deployed by using Cloud Container Engine (CCE), applications created by using ServiceStage, and components deployed on Elastic Cloud Server (ECS) are monitored. For container monitoring, only workloads deployed by using CCE and applications created by using ServiceStage are monitored. For details, see [Component Monitoring](#).

3.3.5 Metric Monitoring

Metric monitoring displays metric data of each resource. You can monitor metric values and trends in real time, add desired metrics to a dashboard, create alarm rules, and export monitoring reports. 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 Select up to 12 metrics to be monitored.

Step 3 Set metric parameters according to [Table 3-6](#), view the metric graphs on the right, and analyze metric data from multiple dimensions.

Table 3-6 Metric parameters

Parameter	Description
Time Range	Time period when metrics are monitored.
Statistical Cycle	Interval at which metric data is collected.
Statistic Method	Method used to measure metrics.

 **NOTE**

The number of samples equals to the count of data points.

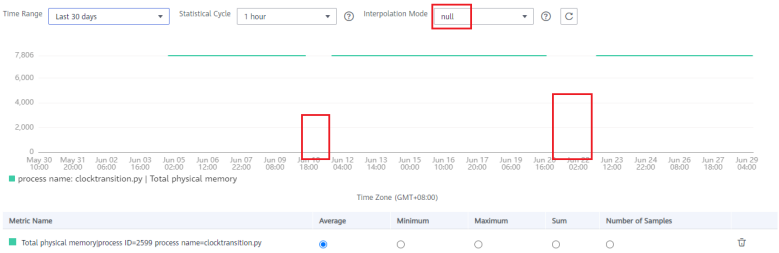
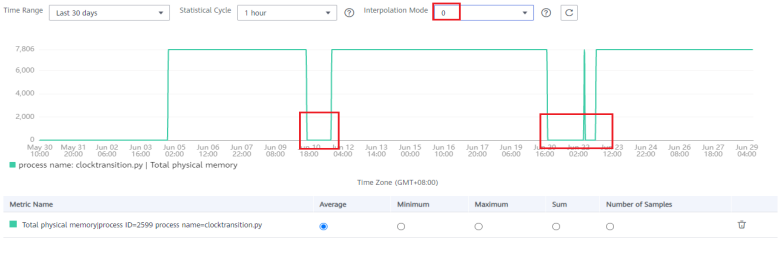
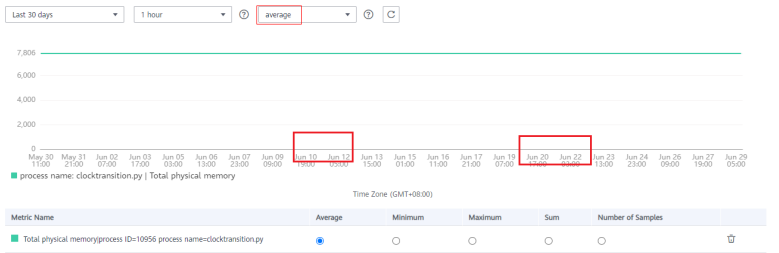
----End

More Operations

You can also perform the operations described in [Table 3-7](#).

Table 3-7 Related operations

Operation	Description
Adding a metric graph to a dashboard	Click Add to Dashboard to add a metric graph to a dashboard.
Adding a threshold rule for a metric	Click Add to Threshold Rule , set threshold rule parameters, and click Submit to add a threshold rule.
Exporting a monitoring report	Click Export Report to export a metric graph as a CSV file to a local PC.

Operation	Description
<p>Setting the interpolation mode</p>	<p>By default, AOM uses null to represent breakpoints in a metric graph, as shown in Figure 3-18. However, a metric graph with breakpoints is not suitable for reports or presentation. To solve the problem, set the value of Interpolation Mode to 0 or average to interpolate values. In this way, you can replace the missing metric data and avoid breakpoints.</p> <p>The value of Interpolation Mode can be null, 0, or average.</p> <ul style="list-style-type: none"> ● null: Breakpoints are represented by null by default, as shown in the following figure. <p>Figure 3-18 Graph when Interpolation Mode is null</p>  <ul style="list-style-type: none"> ● 0: Breakpoints are represented by 0, as shown in the following figure. <p>Figure 3-19 Graph when Interpolation Mode is 0</p>  <ul style="list-style-type: none"> ● average: Breakpoints are represented by average values, as shown in the following figure. <p>Figure 3-20 Graph when Interpolation Mode is average</p> 

Operation	Description
	<p>NOTE</p> <p>If the value of Interpolation Mode is set to average, breakpoints will be represented by average values. The following describes how to calculate average values.</p> <p>A metric graph may have multiple breakpoints. When multiple breakpoints exist, values will be interpolated for these breakpoints from left to right. The following uses the first breakpoint in a graph as an example to describe the method of calculating the average value. This method can also be applied to other breakpoints.</p> <ul style="list-style-type: none"> • If the first breakpoint is at the start of a metric graph, the value of the breakpoint is the first valid data from its next point to the right. For example, if a metric graph has points a, b, c, d, and e, where a = Null, b = Null, c = Null, d = Null, and e = 5, the value of the first breakpoint (that is, point a) is 5. • If the first breakpoint is in the middle of a metric graph, there are the following two scenarios: Scenario 1: If the values of the previous and next points of the breakpoint are valid, the value of the breakpoint is the average value of these two points. For example, if a metric graph has points a, b, c, d, and e, where a = 1, b = Null, c = 3, d = Null, and e = 5, the value of the first breakpoint (that is, point b) is $(a + c)/2 = (1 + 3)/2 = 2$. Scenario 2: If the value of the previous point of the breakpoint is valid and the value of its next point is null, the value of the breakpoint is the average value of its previous point and the first valid data from its next point to the right. For example, if a metric graph has points a, b, c, d, and e, where a = 1, b = Null, c = Null, d = Null, and e = 5, the value of the first breakpoint (that is, point b) is $(a + e)/2 = (1 + 5)/2 = 3$. Because values are interpolated for breakpoints from left to right, the value of the second breakpoint (that is, point c) is $(b + e)/2 = (3 + 5)/2 = 4$, the value of the third breakpoint (that is, point d) is $(c + e)/2 = (4 + 5)/2 = 4.5$. • If the first breakpoint is at the end of a metric graph, the value of the breakpoint is the value of the previous point. For example, if a metric graph has points a, b, c, d, and e, where a = 1, b = 2, c = 3, d = 4, and e = Null, the value of the first breakpoint (that is, point e) is 4. • If all points in a metric graph are breakpoints, the values of all these points are still null, even though you set the value of Interpolation Mode to average. For example, if a metric graph has points a, b, c, d, and e, where a = Null, b = Null, c = Null, d = Null, and e = Null, the values of all breakpoints are null.

3.4 Log Management

3.4.1 Searching for Logs

AOM enables you to quickly query logs, and use log source information and context to locate faults.





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

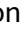
Step 2 On the **Log Search** page, click the **Component**, **System**, or **Host** tab and set filter criteria as prompted.

 **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 case-sensitive keyword.
 - Enter a keyword for exact search. A keyword refers to a word between two adjacent delimiters.
 - Enter a keyword containing an asterisk (*) or a question mark (?) for fuzzy search. For example, enter **ER?OR**, ***ROR**, or **ER*R**.
 - Enter a phrase for exact search. For example, enter **Start to refresh** or **Start-to-refresh**. Note that hyphens (-) are delimiters.
 - Enter a keyword containing AND (&&) or OR (||) for search. For example, enter **query logs&&error*** or **query logs||error**.
 - If no log is found, you are advised to narrow down the search scope and add an asterisk (*) before and after the keyword for fuzzy match.
 - Entering both Chinese and English to search logs is not allowed.

Step 3 View the search results 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 change the order.  indicates the default order.  indicates the ascending order of time (that is, the latest log is displayed at the end of the list).  indicates the descending order of time (that is, the latest log is displayed at the beginning of the list).

1. Click  on the left of the log list to view details.
2. AOM allows you to view the surrounding 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 log 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 more than or equal to 100 logs printed prior to a log and more than or equal to 99 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  on the right of the **Log Search** page, select an export format, and export the search result to a 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 the 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, as shown in **Figure 3-21**. If you select the TXT format, only log content can be exported, as shown in **Figure 3-22**. Each row represents a log. If a log contains a large amount of content, view the log using Notepad++.

Figure 3-21 Exporting logs in CSV format

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
1	Time	Type	Service Name	Instance/Process Name	Host IP Address	Namespace	Cluster Name	Source	Description											
2	2018-12-11	Service	a12345asd	a12345asd	172.16.0.95	default	lycluster1211	/var/ICA/	2018-12-18 16:14:09.089 (5397)[W] ntp_linux.go:36 update ntpStatus: &{status:1 serverStatus:1 offset:}											
3	2018-12-11	Service	a12345asd	a12345asd	172.16.0.95	default	lycluster1211	/var/ICA/	2018-12-18 16:14:09.089 (5397)[W] ntp_linux.go:107 NTPConfig has no set the main NTP_Server!											
4	2018-12-11	Service	a12345asd	a12345asd	172.16.0.95	default	lycluster1211	/var/ICA/	2018-12-18 16:13:58.626 (5397)[W] container_watcher.go:359 get label by pod[evs-driver-fknb6] fail, podName2podInfoM: map[]											
5	2018-12-11	Service	a12345asd	a12345asd	172.16.0.95	default	lycluster1211	/var/ICA/	2018-12-18 16:13:58.626 (5397)[W] container_watcher.go:359 get label by pod[obs-driver-lfhjg] fail, podName2podInfoM: map[]											
6	2018-12-11	Service	a12345asd	a12345asd	172.16.0.95	default	lycluster1211	/var/ICA/	2018-12-18 16:13:58.626 (5397)[W] container_watcher.go:359 get label by pod[sfs-driver-f85hn] fail, podName2podInfoM: map[]											
7	2018-12-11	Service	a12345asd	a12345asd	172.16.0.95	default	lycluster1211	/var/ICA/	2018-12-18 16:13:58.626 (5397)[W] container_watcher.go:359 get label by pod[storage-driver-z5rv2] fail, podName2podInfoM: map[]											
8	2018-12-11	Service	a12345asd	a12345asd	172.16.0.95	default	lycluster1211	/var/ICA/	2018-12-18 16:13:58.626 (5397)[W] container_watcher.go:359 get label by pod[atps-7cc556659b-hvk57] fail, podName2podInfoM: map[]											
9	2018-12-11	Service	a12345asd	a12345asd	172.16.0.95	default	lycluster1211	/var/ICA/	2018-12-18 16:13:58.626 (5397)[W] container_watcher.go:359 get label by pod[atps-7cc556659b-mp8cm] fail, podName2podInfoM: map[]											
10	2018-12-11	Service	a12345asd	a12345asd	172.16.0.95	default	lycluster1211	/var/ICA/	2018-12-18 16:13:58.626 (5397)[W] container_watcher.go:359 get label by pod[atps-7cc556659b-qh47x] fail, podName2podInfoM: map[]											

Figure 3-22 Exporting logs in TXT format

```

2018-12-18 16:14:09.089 (5397) [W] ntp_linux.go:36 update ntpStatus: &{status:1 serverStatus:1 offset:}
2018-12-18 16:14:09.089 (5397) [W] ntp_linux.go:107 NTPConfig has no set the main NTP_Server!
2018-12-18 16:13:58.626 (5397) [W] container_watcher.go:359 get label by pod[evs-driver-fknb6] fail, podName2podInfoM: map[]
2018-12-18 16:13:58.626 (5397) [W] container_watcher.go:359 get label by pod[obs-driver-lfhjg] fail, podName2podInfoM: map[]
2018-12-18 16:13:58.626 (5397) [W] container_watcher.go:359 get label by pod[sfs-driver-f85hn] fail, podName2podInfoM: map[]
2018-12-18 16:13:58.626 (5397) [W] container_watcher.go:359 get label by pod[storage-driver-z5rv2] fail, podName2podInfoM: map[]
2018-12-18 16:13:58.626 (5397) [W] container_watcher.go:359 get label by pod[atps-7cc556659b-hvk57] fail, podName2podInfoM: map[]
2018-12-18 16:13:58.626 (5397) [W] container_watcher.go:359 get label by pod[atps-7cc556659b-mp8cm] fail, podName2podInfoM: map[]
2018-12-18 16:13:58.626 (5397) [W] container_watcher.go:359 get label by pod[atps-7cc556659b-qh47x] fail, podName2podInfoM: map[]
2018-12-18 16:13:58.626 (5397) [W] container_watcher.go:359 get label by pod[atps-7cc556659b-pjdhv] fail, podName2podInfoM: map[]

```

----End

3.4.2 Viewing Log Files

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

- 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 name. Information such as the log file name and latest written time is displayed on the right of the page.
- Step 3** Click **View** in the **Operation** column of the desired instance. **Table 3-8** shows how to view log file details. **Figure 3-24** shows log file details.

Table 3-8 Operations


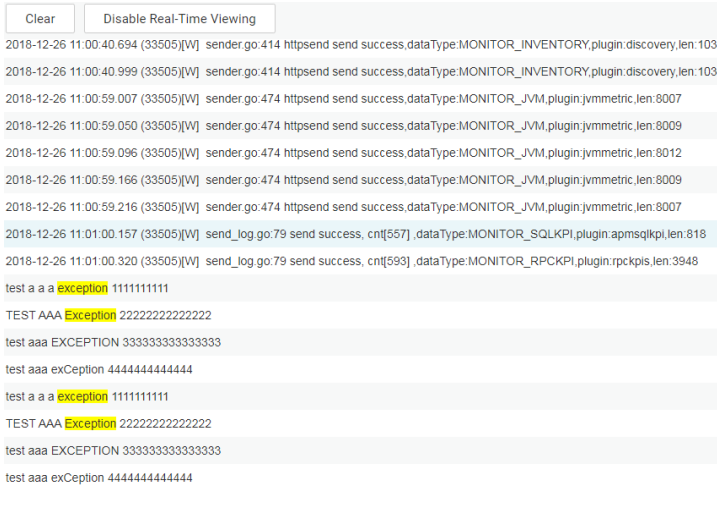


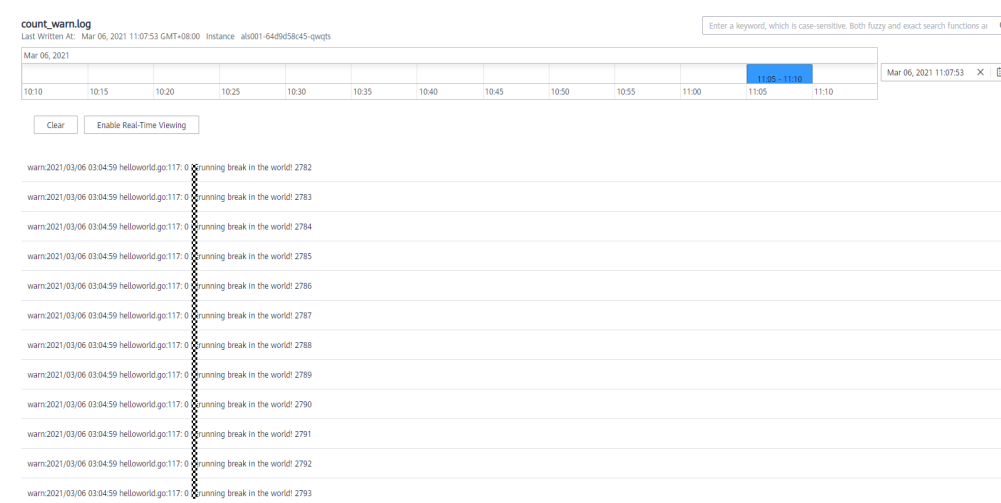
Operation	Setup	Description
Setting a time range	Date and time	Click  to select the date and time.
	Time range	Click the desired time on the time axis to set a time range. You can select only one unit (5 minutes) each time.
Viewing log files	Clear	Click Clear to clear the logs displayed on the screen. This operation clears only the logs displayed on the screen but does not delete them.
	Viewing real-time logs	<p>The function of real-time monitoring is disabled by default. To enable it, click Enable Real-Time Viewing. 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, as shown in the following figure.</p> <p>Figure 3-23 Viewing real-time logs</p> 
Viewing log files	Maximized display	Click  to maximize a page. Components like the time axis are invisible on the screen. Click  again to cancel the maximized display.

Figure 3-24 Log file details



----End

3.4.3 Configuring VM Log Collection Paths

AOM can collect and display VM logs. VM refers to an Elastic Cloud Server (ECS) running Linux. Before collecting logs, configure a log collection path according to the following procedure.

Prerequisites

You have installed the ICAgent on a VM according to [Installing the ICAgent](#). About five minutes after the ICAgent is installed, you can view your VM in the VM list on the **Path Configuration** page.

Precautions

- The ICAgent collects ***.log**, ***.trace**, and ***.out** log files only. For example, **/opt/yilu/work/xig/debug_cpu.log**.
- Ensure that an absolute path of the 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.
- For ECSs in the same resource set, only the latest log collection configuration in the system will be used. AOM and LTS log collection configurations cannot take effect at the same time. 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.

Configuring Log Collection Paths for a Single VM on the Console

- Step 1** Log in to the AOM console. In the navigation pane, choose **Log > Path Configuration** and click the **Host Log** tab.
- 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 files that have file handles and are suffixed with **.log**, **.trace**, or **.out** on the page.


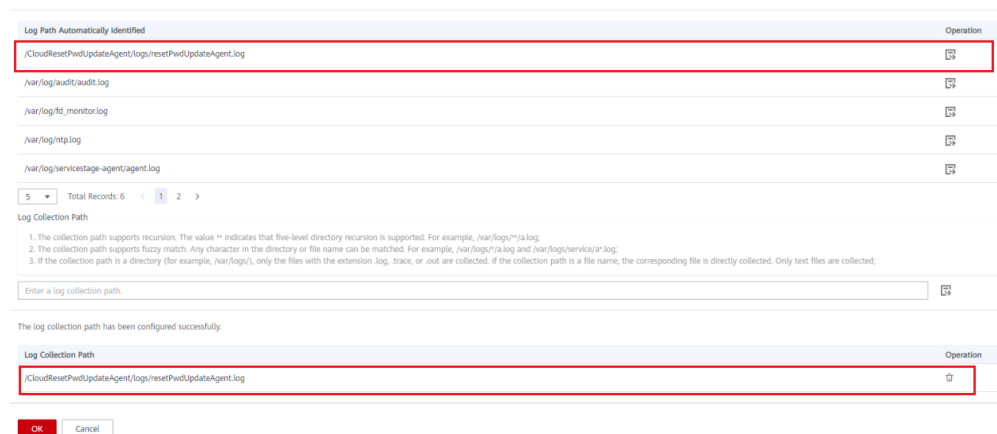
You can click  in the **Operation** column to add a path automatically identified by the ICAgent to the log collection path list. To configure multiple paths, repeat this operation.

Figure 3-25 Using the paths automatically identified by the ICAgent



- **Manually configuring log collection paths**


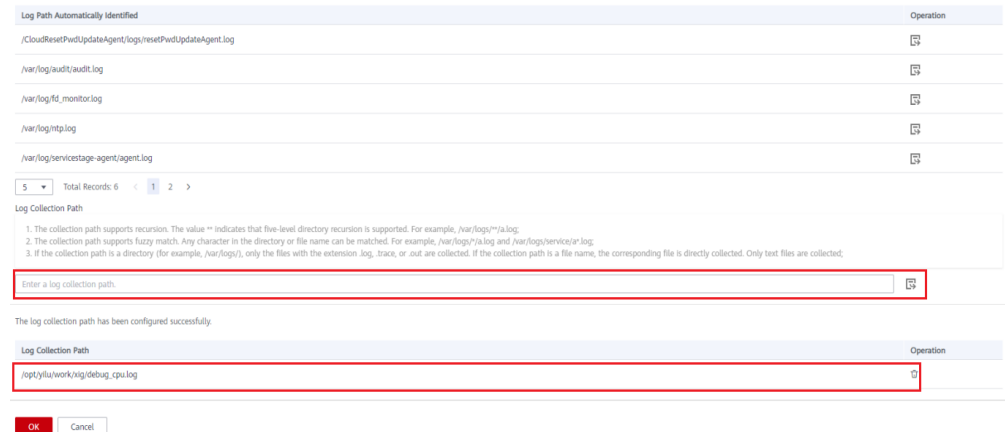
If the paths automatically identified by the ICAgent cannot meet your requirements, enter a log directory or file (for example, **/opt/yilu/work/xig/debug_cpu.log**) in the **Log Collection Path** text box, and then click  to add the path to the log collection path list. To configure multiple paths,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

Figure 3-26 Manually configuring log collection paths



Step 3 Click **OK**.

----End

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

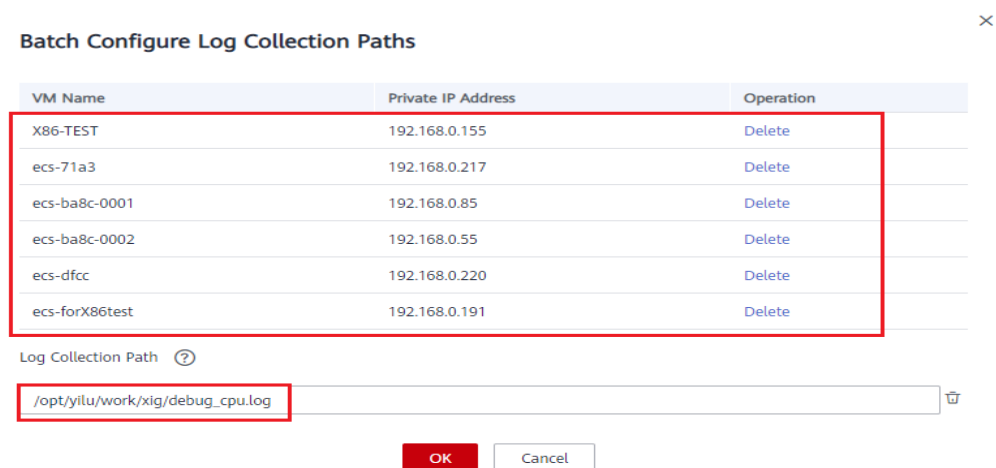
You can configure log collection paths for multiple VMs in batches. When your component is deployed on multiple VMs, configuring log collection paths in batches helps you greatly reduce workload.

Step 1 Log in to the AOM console. In the navigation pane, choose **Log > Path Configuration** and click the **Host Log** tab.

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. To configure multiple paths, click **Add Log Collection Path**.

Figure 3-27 Configuring log collection paths in batches



 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 a log collection path is configured, the ICAGENT collects log files from the configured path. The collection takes about 1 minute. After the collection is complete, 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 [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 [Searching for Logs](#).

3.5 Configuration Management

3.5.1 Agent Management

3.5.1.1 Installing the ICAGENT

The following table describes the ICAGENT status.

Table 3-9 ICAGENT status

Status	Description
Running	The ICAGENT is running properly.
Uninstalled	The ICAGENT is not installed. For details about how to install the ICAGENT, see Installing the 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 Through Logging In to the Server and then install it again.

Status	Description
Upgrading	The ICAgent is being upgraded. This operation takes about 1 minute to complete.
Upgrade failed	Failed to upgrade the ICAgent. Uninstall the ICAgent according to Uninstalling the ICAgent Through Logging In to the Server and then install it again.
Offline	The Access Key ID/Secret Access Key (AK/SK) are incorrect. Obtain the correct AK/SK and install the ICAgent again.
Abnormal	The ICAgent is abnormal. Contact technical support.
Restricted	The AOM license is restricted. Check the license and update it in a timely manner.

Prerequisites

Before installing the 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 the ICAgent. Note that the two methods are not applicable to container nodes created using ServiceStage or Cloud Container Engine (CCE). For container nodes, you do not need to manually install the ICAgent. Instead, you only need to perform certain operations when creating clusters or deploying applications.

For details, see [Table 3-10](#).

Table 3-10 Installation methods

Method	Application Scenario
Initial installation	This method is used when the following conditions are met: <ol style="list-style-type: none"> 1. An elastic IP address (EIP) has been bound to the server. 2. The ICAgent has never been installed on the server.
Inherited installation	This method is used when the following conditions are met: You have multiple servers on which the ICAgent is to be installed. One server is bound to an EIP, but others are not bound to an EIP. You can use this method to install the ICAgent on the servers that are not bound to an EIP.

Initial Installation

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

Step 1 Obtain an AK/SK.

- If you have obtained the AK/SK, skip this step.
- If you have not obtained the AK/SK, obtain them first.

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

Step 3 Click **Install ICAgent**.

Step 4 Generate and copy the ICAgent installation command.

1. Enter the obtained AK/SK in the text box to generate the ICAgent installation command.

 **NOTE**

Ensure that the AK/SK are correct. Otherwise, the ICAgent cannot be installed.

2. Click **Copy Command**.

Step 5 Use a remote login tool, such as PuTTY, to log in as user **root** to the server where the ICAgent is to be installed and run the command copied in the previous step to install the ICAgent.

Step 6 Run the command copied in **Step 4** and enter the obtained AK/SK as prompted to install the ICAgent.

----End

 **NOTE**

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

Inherited Installation

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

NOTICE

After the ICAgent is upgraded, the **/opt/ICAgent/** directory and the files stored in it will be deleted. Therefore, reinstall the ICAgent and then perform inherited installation.

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/remotelInstall/remote_install.sh -ip  
x.x.x.x
```

Step 2 Enter the password of the **root** user of the server where the ICAgent is to be installed as prompted.

 **NOTE**

- If both the Expect tool and the ICAgent have been installed on the server, the ICAgent will be installed on 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.
- 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 installed.
- Ensure that the **ICProbeAgent.tar.gz** installation package is transmitted to the server to be installed.
- If the message **ICAgent install success** is displayed, the ICAgent is successfully installed in the **/opt/oss/servicemgr/** directory. After the ICAgent is successfully installed, choose **Configuration Management > Agent Management** to view the ICAgent status.
- If the ICAgent fails to be installed, uninstall it according to [Uninstalling the ICAgent Through Logging In to the Server](#) and then install it again. If the problem persists, contact technical support.

----End

Inherited Batch Installation

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 install the ICAgent on multiple remote servers in batches with a few clicks.

NOTICE

1. Ensure that you can run the **SSH** and **SCP** commands on the ECS server where the ICAgent has been installed to communicate with the remote ECS servers where the ICAgent is to be installed.
2. If you have installed the ICAgent in a server through an agency, you also need to set an agency for other servers where the ICAgent is to be installed.
3. Batch installation scripts depend on Python versions. You are advised to implement batch installation on hosts running Python 2.x. Python 3.x does not support batch installation.
4. You need to press **Enter** at the end of each line in the **iplist.cfg** file.
5. After the ICAgent is upgraded, the **/opt/ICAgent/** directory and the files stored in it will be deleted. Therefore, reinstall the ICAgent and then perform inherited batch installation.

Prerequisites

The IP addresses and passwords of all servers for which the 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 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 after use.
- If the passwords of all servers are the same, only list IP addresses in the **iplist.cfg** file and enter the password once during execution. If the password of an IP address is different from those of the other ones, list both passwords and IP addresses in the **iplist.cfg** file.
- The batch installation function depends on Python 2.7.*. If the system displays a message indicating that Python cannot be found during the installation, install Python 2.7.* and try again.

Procedure

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

```
bash /opt/oss/servicemgr/ICAgent/bin/remotelInstall/remote_install.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 install begin  
start to install python pexpect module  
use local pexpect package  
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 is successfully installed on all the hosts listed in the configuration file.

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

----End

3.5.1.2 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: user-defined nodes** 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 a time. If you select **Other: user-defined nodes** in **Step 2**, select a desired host and then click **Upgrade ICAgent**.
- Step 4** Wait for about 1 minute to complete the upgrade. When the ICAgent status changes from **Updating** to **Running**, the ICAgent is successfully upgraded.

 **NOTE**

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

----End

3.5.1.3 Uninstalling the ICAgent

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

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

- **Uninstalling the ICAgent Through the AOM Console:** Applies to the scenario where the ICAgent has been successfully installed and needs to be uninstalled.
- **Uninstalling the ICAgent Through Logging In to the Server:** Applies to the scenario where the ICAgent fails to be installed and needs to be uninstalled for reinstallation.
- **Remotely Uninstalling the ICAgent:** Applies to the scenario where the ICAgent has been successfully installed and needs to be remotely uninstalled.
- **Uninstalling the ICAgent in Batches:** Applies to the scenario where the ICAgent has been successfully installed and needs to be uninstalled in batches.

Uninstalling the ICAgent Through the AOM Console

- Step 1** In the navigation pane, choose **Configuration Management > Agent Management**.
- Step 2** Select **Other: user-defined nodes** 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 **Yes**.

The ICAgent begins to be uninstalled. This operation takes about 1 minute to complete. If the involved server is removed from the node list, the ICAgent is successfully uninstalled.

----End

Uninstalling the ICAgent Through Logging In to the Server

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

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 is successfully uninstalled.

----End

Remotely Uninstalling 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 of the server where the ICAgent is to be uninstalled 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.
- 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 is successfully uninstalled. After the ICAgent is successfully 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 after use.
- If the passwords of all servers are the same, only list IP addresses in the **iplist.cfg** file and enter the password once during execution. If the password of an IP address is different from those of the other ones, list both passwords and IP addresses in the **iplist.cfg** file.
- 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 is successfully uninstalled from all the hosts listed in the configuration file.

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

----End

3.5.2 Configuring Application Discovery

AOM can discover applications and collect their metrics based on configured rules. Application discovery supports both automatic and manual configuration. This section focuses on manual configuration.

- **Automatic configuration**

After you install the ICAgent on a host according to [Installing the ICAgent](#), the ICAgent automatically discovers applications on the host based on **Built-**

in [Service Discovery Rules](#) and displays them on the **Application Monitoring** page.

- **Manual configuration**

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

Filtering Rules

The ICAgent will periodically detect processes on the target host. The effect is similar to that of running the `ps -e -o pid,comm,lstart,cmd | grep -v defunct` command. Then, the ICAgent checks whether processes match the filtering rules in [Table 3-11](#). 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 and is discovered by AOM.

ICAgent detection results may as follows:

```
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 3-11 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) .

Filtering Rule	Example
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 Service Discovery Rules

AOM provides **Default_Rule**, ServiceStage provides **servicestage-default-rule**, and ICAgent has the built-in **Sys_Rule**. These rules are executed on all hosts, including those added later. The priorities of the three built-in discovery rules are as follows: **Sys_Rule** > **servicestage-default-rule** > **Default_Rule**. Rule details are as follows:

Sys_Rule (cannot be disabled)

- For the component name, obtain the value of **-Dapm_tier** in the command, the value of the environment variable **PAAS_APP_NAME**, and the value of **-Dapm_tier** of the environment variable **JAVA_TOOL_OPTIONS** based on the priorities in descending order.
- For the application name, obtain the value of **-Dapm_application** in the command, the value of environment variable **PAAS_MONITORING_GROUP**, and the value of **-Dapm_application** in the environment variable **JAVA_TOOL_OPTIONS** based on the priorities in descending order.

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
```

NOTE

sys_rule (built-in application discovery rule) is used to discover CCE workloads in the CCE scenario (except the APM scenario).

servicestage-default-rule (can be disabled)

Detect the processes whose environment variables contain **CAS_APPLICATION_NAME**, **CAS_COMPONENT_NAME**, and **CAS_ENVIRONMENT_NAME**.

Obtain the value of the environment variable **CAS_COMPONENT_NAME** and combine it as the component name.

Obtain the value of the environment variable **CAS_APPLICATION_NAME** and combine it as the application name.

 NOTE

Nginx of a version earlier than 1.19 does not support manual configuration of environment variables. Therefore, the data source of the Nginx component configured on ServiceStage may be displayed as **CCE** on the **Component Monitoring** page of AOM. To solve this problem, use Nginx of 1.19 or a later version, manually configure environment variables such as **CAS_COMPONENT_NAME** and **CAS_COMPONENT_ID**.

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

The priorities of discovery rules are as follows: Sys_Rule > Custom discovery rules > Default_Rule. **servicestage-default-rule** is also a custom discovery rule.

Step 1 In the navigation pane, choose **Configuration Management > Service 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, **ruletest**.
2. Select a typical host, for example, **hhhhh-27465**, 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 need to add one check item at least and can add five check items at most. 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, go to the next step.

Step 5 Set a component name.

1. Set an application name.


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

 **NOTE**

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

In the **Component Name Settings** area, click **Add Naming Rule** to set a component name for the discovered process.

 **NOTE**

- The component name cannot be left blank.
 - When you add multiple naming rules, all the naming rules are combined as the component name of the process. Metrics with the same component name are aggregated.
3. Preview the component name.
- If the application or component name does not meet your requirements, click  in the **Preview Component Name** table for renaming.

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.

 **NOTE**

- Do not use multiple custom discovery rules with the same priority for the same process.
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 those added later.

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

Step 8 Wait for about two minutes, choose **Monitoring > Component Monitoring** in the navigation pane, select the **hhhhh-27465** host from the cluster drop-down list, and find out the **/openswitch/** component that has been monitored.

----End

 **NOTE**

Custom discovery rules cannot be used to discover workload processes in the CCE scenario. You are advised to use custom discovery rules to discover non-workload processes in the CCE scenario or processes in the VM scenario.

More Operations

After creating an application discovery rule, you can also perform the operations described in [Table 3-12](#).

Table 3-12 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 process metrics.
Deleting a rule	<ul style="list-style-type: none"> To delete a discovery rule, click Delete in the Operation column. <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>

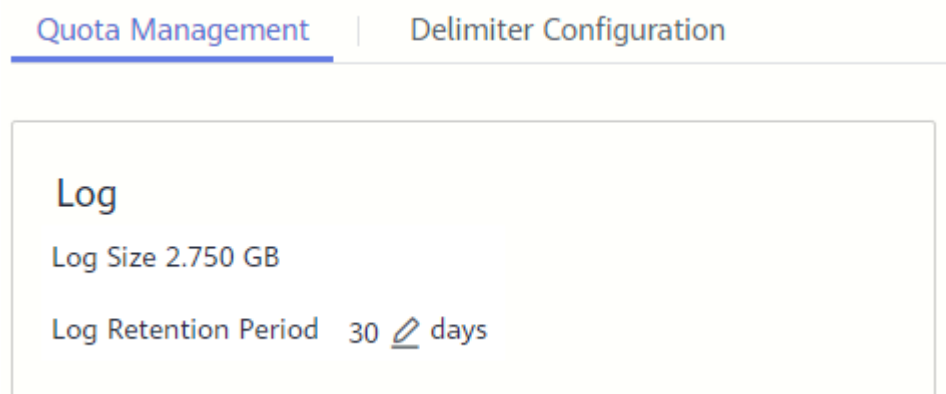
3.5.3 Log Configuration

3.5.3.1 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.

Figure 3-28 Viewing logs



Log retention period: 30 days.

----End

3.5.3.2 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 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.

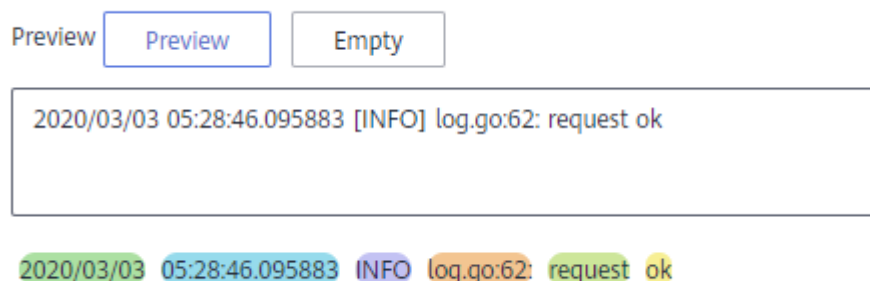
You can configure delimiters using the following methods: If you use both methods at the same time, the union set will be selected.

- Custom delimiters: Click , enter a delimiter in the text box, and click .
- Use ASCII code: 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 preview in the text box and click **Preview**.

Figure 3-29 Previewing logs



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 3-13 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
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

ASCII Value	Control Character	ASCII Value	Control Character	ASCII Value	Control Character	ASCII Value	Control Character
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
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

ASCII Value	Control Character	ASCII Value	Control Character	ASCII Value	Control Character	ASCII Value	Control Character
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)

4 FAQs

4.1 What Can I Do If an ICAgent Is Offline?

After an ICAgent is installed, its status is offline.



Node Name	Node IP Address	ICAgent Status	ICAgent Version	Java Probe Version	Updated At
ecs-63a2-CLX1	192.168.0.62	Running	5.12.199	6.0.0	Mar 11, 2022 09:45:45 GMT+08:00
ecs-cassandra	192.168.0.3	Running	5.12.199	1.0.48	Mar 11, 2022 10:03:34 GMT+08:00
ecs-f790-CLX	192.168.0.88	Running	5.12.199	6.0.0	Mar 11, 2022 09:45:45 GMT+08:00
ecs-zixzix	192.168.0.178	Offline Details	--	--	--

Problem Analysis

- **Cause:** The AK/SK configuration is incorrect or ports 30200, 30201, 8149, 8923, and 8102 are not connected.
- **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 the command output is displayed, the AK/SK configuration is correct. Go to [Step 4](#).

Step 3 After configuring the AK/SK, reinstall the ICAgent. For details, see [Installing 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 whether ports 30200, 30201, 8149, 8923, and 8102 are connected:

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

- 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 ICAgent. If the installation still fails, contact technical support.

----End

4.2 Obtaining an AK/SK

NOTE

Each user can create a maximum of two 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**. In the dialog box that is displayed, enter the login password and verification code sent to your email or mobile phone.
4. Click **OK** and download the generated AK/SK pair.

NOTE

Keep the AK/SK secure.

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

In 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 want to query metrics in the last hour, the available statistical cycles are 1 minute and 5 minutes.

The following table shows the relationship between the time range and statistical cycle.

Table 4-1 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, 15 minutes, or 1 hour
Last 1 day	
Last 1 week	15 minutes, 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
Last 30 days	NOTE 1 day is only for the metrics generated based on log statistical rules.

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

Resource statuses include **Normal**, **Warning**, **Abnormal**, **Deleted**, and **Silent**. **Warning**, **Abnormal**, or **Silent** may result in resource running exceptions. You can analyze and rectify faults according to the following guidance.

Warning

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

Suggestion: Handle the alarm based on alarm details.

Abnormal

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

Suggestion: Handle the alarm 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 uninstalled or abnormal. For details on how to solve the problem, see [Table 4-2](#).

Table 4-2 ICAgent troubleshooting suggestions

Status	Suggestion
Uninstalled	Install the ICAgent according to Installing the ICAgent .
Installing	Wait for about 1 minute to complete the ICAgent installation.
Installation failed	Uninstall the ICAgent according to Uninstalling the ICAgent Through Logging In to the Server . Install the ICAgent again.
Upgrading	Wait for about 1 minute to complete the ICAgent upgrade.
Upgrade failed	Uninstall the ICAgent according to Uninstalling the ICAgent Through Logging In to the Server and then install it again.
Offline	Ensure that the Access Key ID/Secret Access Key (AK/SK) or Elastic Cloud Server (ECS) agency configuration is correct.
Faulty	Contact technical support.

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

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

- **Cause 3: The local time of the host is not synchronized with the NTP server time.**

 **NOTE**

NTP Sync Status: indicates whether the local time of the host is synchronized with the NTP server time. The value can be **0** or **1**. **0** indicates the synchronized status while **1** indicates the asynchronized status.

Suggestion: Choose **Monitoring > Metric Monitoring** and check the **NTP Sync Status** metric of the host. If the value of **NTP Sync Status** is **1**, the local time of the host is not synchronized with that of the NTP server. To solve the problem, perform synchronization.


- **Cause 4: The resource is deleted or stopped.**

Suggestions:

- On the ECS page, check whether the host is restarted, stopped, or deleted.
- On the Cloud Container Engine (CCE) page, check whether the component is stopped or deleted.
- If a discovery rule is stopped or deleted, the component discovered based on the rule will also be stopped or deleted. On the AOM page, check whether the discovery rule is stopped or deleted.

4.5 How Can I Do If I Do Not Have the Permission to Access SMN?

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



Problem Analysis

- **Cause:** The IAM user does not have the permission to access Simple Message Notification (SMN).
- **Impact:** Email or message notifications cannot be received.

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.

4.6 How Do I Distinguish Alarms and Events?

Similarities Between Alarms and Events

Alarms and events refer to the information reported to Application Operation Management (AOM) when the status of AOM or an external service, such as ServiceStage, Cloud Container Engine (CCE), or Application Performance Management (APM) changes.

Differences Between Alarms and Events

- Alarms are reported when AOM or an external service, such as ServiceStage, CCE, or APM is abnormal or may cause exceptions. Alarms need to 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, CCE, or APM encounters some changes. Such changes do not necessarily cause service exceptions. Events do not need to be handled.

4.7 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.

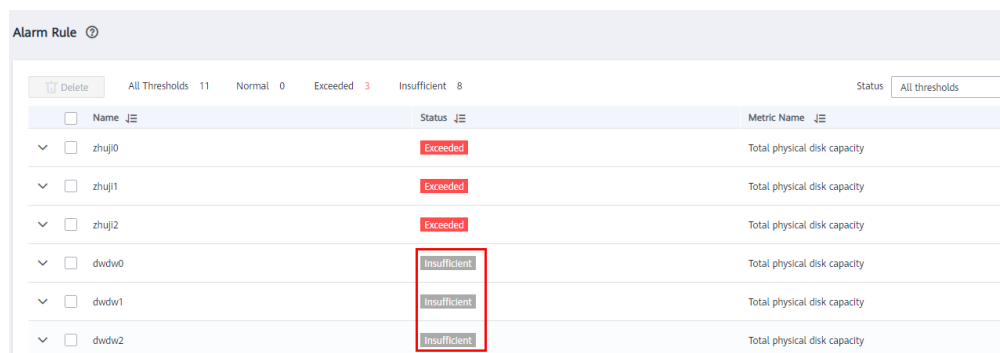
There is a time interval between log collection and processing. If the number of logs is small, the latency is about 10s. If the number of logs is large, the latency is much longer.

4.8 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.

4.9 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, as shown in the following figure.



Name	Status	Metric Name
zhuj0	Exceeded	Total physical disk capacity
zhuj1	Exceeded	Total physical disk capacity
zhuj2	Exceeded	Total physical disk capacity
dwdw0	Insufficient	Total physical disk capacity
dwdw1	Insufficient	Total physical disk capacity
dwdw2	Insufficient	Total physical disk capacity

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 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.

4.10 Why the Status of a Workload that Runs Normally Is Abnormal on the AOM Page?

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

Workload	Status	Cluster	Namespace
coredns	Abnormal	cluster-factory	kube-system
storage-driver	Abnormal	cluster-factory	kube-system

Possible causes:

- a. The ICAgent version is too early.

Currently, the ICAgent needs to be upgraded by users. However, if the ICAgent version is too early, 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.

<input type="checkbox"/>	Node Name	Node IP Address	ICAgent Status	ICAgent Version
<input type="checkbox"/>	Dont-Touch	192.168.0.67	Running	5.13.53
<input type="checkbox"/>	as-config-ljib-UD855PVU	192.168.0.26	Running	5.13.53

- b. 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.