Application Operations Management

Service Overview

Issue	01
Date	2026-01-31

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: https://www.huaweicloud.com/intl/en-us/

Contents

1 AOM Infographic

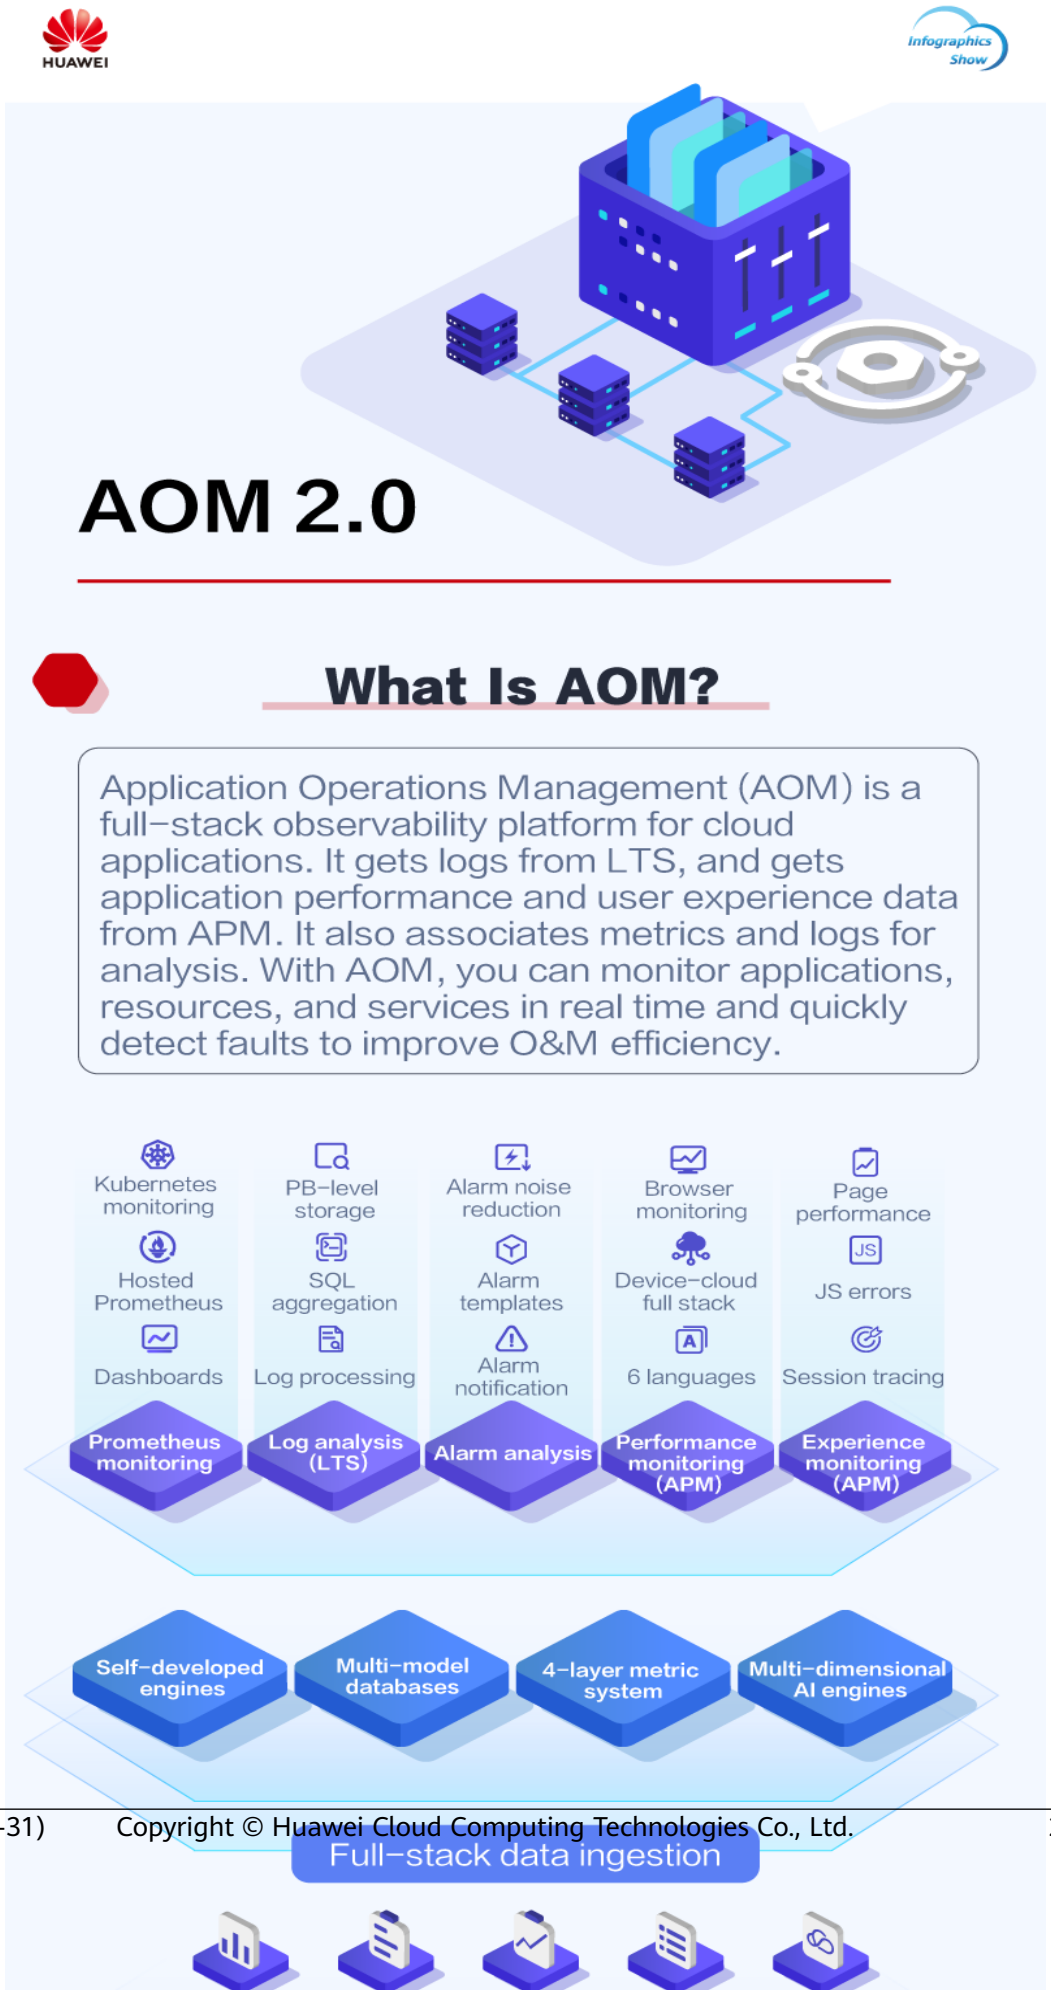

2 What Is AOM?

Application Operations Management (AOM) is a full-stack observability platform for cloud applications. It integrates observability data from Cloud Eye, Log Tank Service (LTS), and Application Performance Management (APM) and associates metrics, logs, and traces for analysis. With AOM, you can monitor applications, resources, and services in real time and quickly detect faults to improve O&M efficiency.

Figure 2-1 AOM panorama

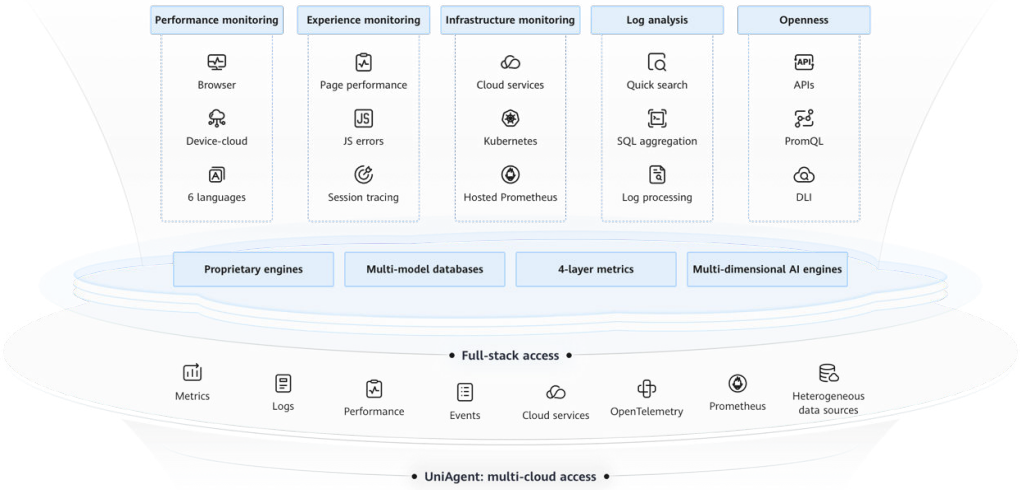

- **Access Center**: AOM provides all-scenario data ingestion capabilities. Metrics, traces, and logs can be ingested from the business layer (Android/iOS/web), component layer (Java/Go), middleware layer (Redis/MySQL), running environment layer (ECS/Kubernetes), and cloud service layer, and through APIs and open-source protocols (OpenTelemetry/Kafka). Also, UniAgent is used to manage the lifecycle of data collection plug-ins.
- **Alarm Management**: AOM provides unified Prometheus/log/cloud service alarm rule, template, and notification object management. Over 100 preset alarm templates are available. Alarm grouping, suppression, and silence are supported to reduce alarm noise. AOM also supports custom message templates and various notification channels (emails/SMS/voice calls/WeCom/DingTalk/Lark/Webhook).

- **Dashboard**: AOM supports custom dashboards and various graphs such as tables, line graphs, pies, and maps. There are more than 30 preset dashboards ready to use. You can add custom variables and filters to filter metric and log data for operations analysis.
- **Prometheus Monitoring**: AOM is fully interconnected with the open-source Prometheus ecosystem. It provides fully hosted Prometheus services, and supports PromQL-based queries, remote reads, remote writes, metric browsing, and interconnection with Grafana.
- **Infrastructure Monitoring**: AOM collects monitoring data from hosts (processes) and containers (workloads/pods), displays infrastructure data (such as the host CPU/memory/disk usage and logs), and associates data for analysis.
- **Log Management**: AOM provides all-scenario log collection, search of tens of billions of logs in seconds, PB-level storage, one-stop log processing, visualized log SQL analysis, and log transfer capabilities, meeting application O&M, security compliance, and operations analysis requirements.
- **Application Performance & User Experience Monitoring**: AOM analyzes the latency and exceptions from webs/apps to application microservices/ databases/middleware, enabling O&M personnel to quickly locate root causes and identify application performance bottlenecks, thereby protecting user experience.
- **Cloud Service Monitoring**: AOM monitors cloud service resources such as Elastic Cloud Servers (ECSs) and bandwidth, enabling you to understand cloud resource usage and service running and make fast response to exceptions.

3 Advantages

- **Compatibility and openness**

  AOM supports various open-source protocols, opens O&M data query APIs and collection standards, and provides fully hosted, O&M-free, and cost-efficient cloud native monitoring capabilities.

- **Ready-to-use**

  You can connect applications to AOM without changing code. Data can be collected in a non-intrusive way.

- **Abundant data sources**

  AOM integrates multiple types of data (such as cloud monitoring, logs, application performance, real user experience, and backend connections) for observability analysis.

- **Full-stack integrated monitoring**

  AOM monitors data of clients, servers, and cloud products. It supports data discovery and display, and reports alarms when there are exceptions. It implements integrated monitoring from top to bottom and from the frontend to the backend.

- **Association analysis**

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

- **Precise alarm reporting**

  AOM has a unified alarm system, covering metric, event, and log alarms. It provides alarm noise reduction policies, such as grouping, suppression, and silence. It also supports alarm notification and subscription, so that you can easily cope with alarm storms and detect and clear alarms.

- **Unified visualization**

  Multiple data sources can be monitored and analyzed in the same dashboard. They are displayed in various graphs (such as line and digit graphs), helping you better monitor resources, learn about trends, and make decisions.

4 Scenarios

Scenario 1: Container Monitoring

Pain Points

Prometheus is an ideal monitoring tool for containers. However, self-built Prometheus is costly for small- and medium-sized enterprises with limited O&M labor. For large enterprises that grow rapidly, performance bottlenecks are common. Therefore, many enterprises are turning to hosted Prometheus.

Solution

AOM is compatible with the open-source Prometheus ecosystem. It provides fully hosted Prometheus instances and works with Cloud Container Engine (CCE) to provide **Monitoring Center**. AOM monitors applications and resources in real time, collects metrics and events to analyze application health status, and provides comprehensive, clear, and multi-dimensional data visualization capabilities. It is compatible with mainstream open-source components and enables quick fault locating. This solution has the following advantages:

- Monitoring Center is deeply integrated with Prometheus, a mature monitoring project of the Cloud Native Computing Foundation (CNCF). It brings in observability for your cloud native applications by collecting, storing, and presenting O&M data, such as key metrics and events.
- Monitoring Center provides full-stack monitoring from cloud native infrastructure resources to application workloads, enabling you to clearly perceive the infrastructure and application workload status anytime anywhere.
- Monitoring Center monitors Kubernetes clusters, nodes, and pods, enables end-to-end tracing and visualization for services, and provides cluster health diagnosis, greatly speeding up fault analysis and locating.
- Monitoring Center provides ready-to-use add-ons, data collection, and cloud native cluster monitoring. Compared with monitoring developed based on open-source components, it is more competitive in reliability, availability, and deployment.
- Monitoring Center provides lightweight add-ons for metric collection. Compared with the community's Prometheus, it greatly reduces resource usage and facilitates deployment.

Figure 4-1 Monitoring Center architecture

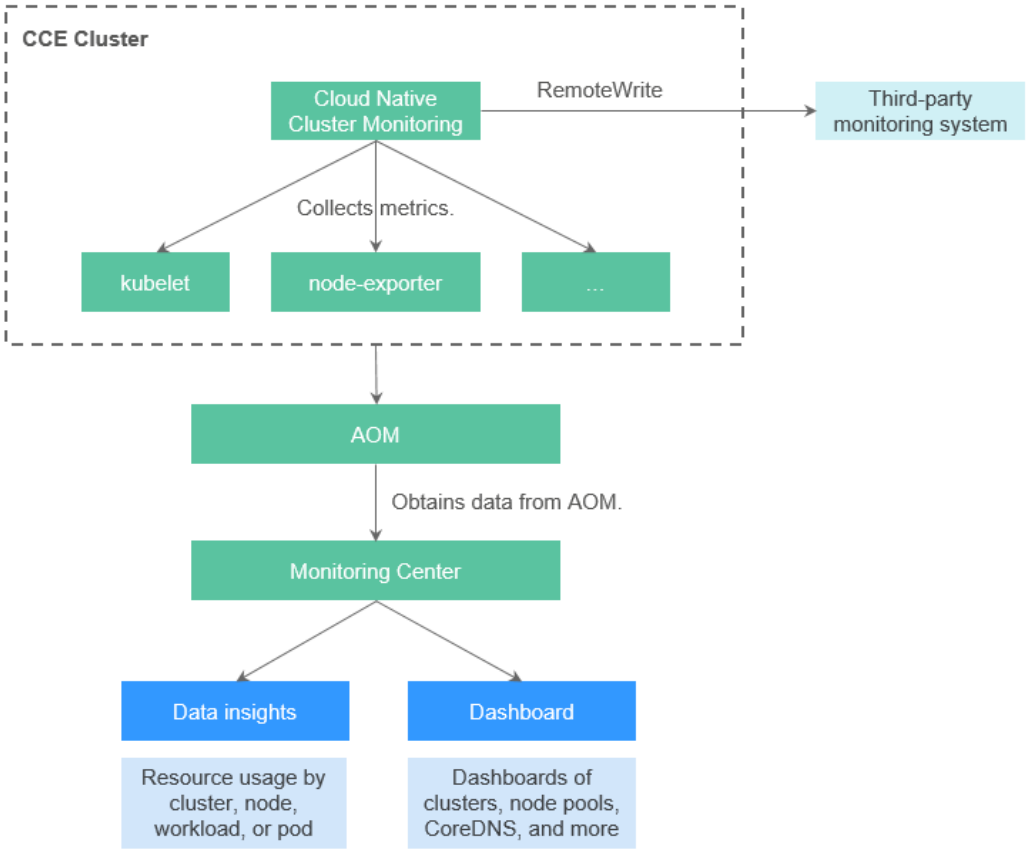

Scenario 2: Application O&M

Pain Points

In the distributed microservice architecture, enterprises can develop diverse applications efficiently, but face great challenges in traditional O&M and diagnosis. For example, an e-commerce application may face:

- Too many alarms

When there are a large number of alarms, it is hard to identify the critical ones. O&M personnel cannot quickly identify big problems and handle them in a timely manner to ensure service running.

- Time-consuming fault locating

  After receiving customers' feedback, customer service personnel submit problems to technical personnel for troubleshooting. In the distributed microservice architecture, a request usually undergoes multiple services/nodes before a result is returned. If a fault occurs, O&M personnel have to view logs on multiple hosts to locate the fault. Even a simple problem can involve multiple teams.

- Lack of unified O&M policies

There are more than 100 service systems. O&M data (logs, metrics, and traces) is scattered on different services, infrastructure, and cloud services. There is no platform for unified data collection, search & analysis, or alarm reporting.

- Complex architecture

  When service logic becomes complex, it is difficult to find out the dependency downstream services (databases, HTTP APIs, and caches) of an application, and external services that depend on the application from the perspective of code. It is also difficult to sort out the service logic, manage the architecture, and plan capacities. For example, enterprises find it hard to determine the number of hosts required in their activities.

Solution

AOM works with Cloud Eye, Application Performance Management (APM), and Log Tank Service (LTS) to provide associated analysis of metrics, traces, and logs. Metrics (such as CPU usage and response time) reflect system performance in a quantitative manner. They help O&M personnel quickly learn about the overall running status of the system. By monitoring and analyzing metrics in real time, O&M personnel can detect performance exceptions in a timely manner. Tracing focuses on the path of a request transferred between components in the system and clearly displays the data flow direction and service call relationships. When a performance metric is abnormal, tracing enables O&M personnel to accurately locate the specific service or component and find out the root cause. Logs record detailed event information during system running, including error stacks and service operation details. After locating a faulty component through tracing, you can view the corresponding logs to obtain the context and further analyze the cause to solve the problem. Metrics, traces, and logs complement each other. Metrics are used for macro monitoring, traces for locating problems, and logs for in-depth analysis of problems. They help build a highly observable system environment.

- Unified data ingestion: AOM provides all-scenario data ingestion capabilities. Metrics, traces, and logs can be ingested from the business layer (Android/iOS/web), component layer (Java/Go), middleware layer (Redis/MySQL), running environment layer (ECS/Kubernetes), and cloud service layer, and through APIs and open-source protocols (OpenTelemetry/Kafka).
- Unified alarms: AOM provides unified metric/log alarm rule, template, and notification object management. Over 100 preset alarm templates are available. Alarm grouping, suppression, and silence are supported to reduce alarm noise. AOM also supports custom message templates and various notification channels (emails/SMS/voice calls/WeCom/DingTalk/Lark/Webhook).
- Dashboard: AOM supports custom dashboards and various graphs such as tables, line graphs, pies, and maps. There are more than 30 preset dashboards ready to use. You can add custom variables and filters to filter metric and log data for operations analysis.
- Topologies and tracing: Application topologies are automatically discovered. Abnormal application instances can be easily detected. After an abnormal application is detected, you can drill down the traces to locate the cause in code.
- SQL analysis: Key metrics (such as number of SQL statement calls, latency, and number of errors) are displayed in graphs. The database performance problems caused by abnormal SQL statements can be analyzed.
- Search for tens of billions of logs in seconds: A high-performance log search engine is provided and a storage-compute decoupled architecture is used.

Multiple tenants can share countless elastic compute resources, enabling log search in seconds.

Scenario 3: Security Compliance

Pain Points

In large enterprises, each department has an independent cloud account for isolating resources. Their O&M personnel rely on logs and alarms to locate and analyze faults. Also, the security department needs to centrally monitor the logs of all departments. Therefore, unified log management of multiple accounts is challenging.

- **Independent O&M by service:** Each service module has an independent account for resource isolation and needs a log service to configure monitoring alarms to quickly locate faults and root causes.
- **Unified log monitoring:** To meet regulatory requirements, the security department needs to aggregate the logs of different department accounts to one account and store them for more than 180 days.

Solution

AOM works with LTS to implement:

- **Independent account management:** Each account has isolated resources and permissions and independently collects its own application and cloud service logs. You can configure alarm rules to demarcate 90% of problems in 10 minutes.
- **Log aggregation across accounts:** Different accounts' logs are copied to one account and stored for at least 180 days for compliance audits, meeting cyber security regulations.

The following shows the solution for centralized log collection and storage beyond 180 days. It complies with the *Cybersecurity Law* and *General Data Protection Regulation (GDPR)*.

Figure 4-2 Security compliance solution

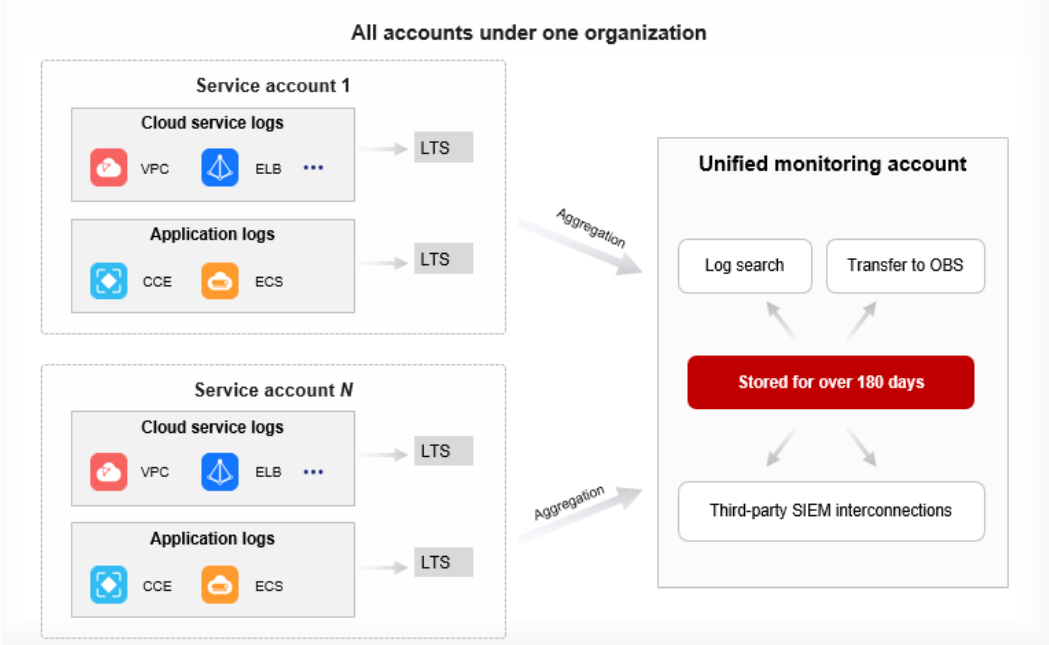

Scenario 4: Operations Analysis

Pain Points

Enterprises collect various logs or metrics (such as those of mobile devices) during their daily operations. After being normalized, filtered, anonymized, and enriched, the logs or metrics can be analyzed using big data platforms and BI tools to provide operations data such as the PV, UV, user stay duration, and transaction amount. Such data helps enterprises understand their operations status, analyze user behavior characteristics, make adjustments in real time, and improve user experience and operations efficiency for digital transformation.

Enterprises often encounter the following pain points during service analysis:

- **Difficult data collection**: The logs of mobile devices (such as web pages, iOS, Android, Baidu/WeChat/DingTalk applets, and quick apps) are hard to collect.
- **Unreliable data transmission:** Mobile device logs/metrics are numerous and frequently transmitted. The transmission is slow and logs are prone to be lost, affecting service analysis.
- **Inconvenient data processing:** Raw data cannot be directly processed by big data platforms.

Solution

AOM works with APM and LTS to collect mobile device logs/metrics and uses the big data platform to analyze service operations. It supports:

- **Full collection of mobile device logs:** SDKs are provided to support functions such as cached data sending, retry upon exceptions, and batch sending. With these SDKs, you can collect mobile device logs or metrics quickly.
- **Fast and reliable reporting:** The collected device logs or metrics are reported in seconds through the transmission link. Moreover, no data will be lost in this process.

- **Visualized dashboards**: You can customize your own dashboards and add various graphs (such as tables, line graphs, pies, and maps) to them. You can also add custom variables and filters to filter data for operations analysis.
- **Data openness to big data platforms**: You can dump metrics to Distributed Message Service (DMS) for Kafka and dump logs to Object Storage Service (OBS), Data Lake Insight (DLI), and Data Warehouse Service (DWS). After the data is processed on the big data platforms, connect to the BI system for visualized display.

This solution can be used to structure logs, analyze logs with SQL syntax, analyze metrics with PromQL syntax, generate visual graphs, and work with big data platforms to help enterprises further explore data value and achieve digital transformation.

Scenario 5: User Experience Improvement

Pain Points

In the Internet era where user experience is of crucial importance, enterprises cannot obtain user access information even if backend services run stably. It is much more difficult to locate frontend problems that occur occasionally. Enterprises suffer higher churn rate if they cannot obtain error details in time.

Solution

AOM works with APM to analyze the complete process (user request > server > database > server > user request) of application transactions in real time, enabling you to monitor comprehensive user experience in real time. For transactions with poor experience, AOM locates problems through topology and tracing.

- Web monitoring monitors page performance, JS error requests, API requests, and service operations metrics (such as PV and UV) in real time.
- Session tracing locates slow requests, loading, and interactions that affect user experience and monitors user usage in real time.
- Page loading performance analysis provides metrics such as First Contentful Paint (FCP), First Paint (FP), and Time to Interactive (TTI), helping you restore user experience and locate the causes of slow access.

5 Functions

This section describes the main functions of Application Operations Management (AOM). You can check if a certain function is available in a region on the console.

Access Center

AOM provides all-scenario data ingestion capabilities. Metrics, traces, and logs can be ingested from the business layer (such as Android, iOS, and webs), component layer (such as Java and Go), middleware layer (such as Redis and MySQL), running environment layer (such as ECS and Kubernetes), and cloud service layer, or using APIs and open-source protocols (such as OpenTelemetry and Kafka). It also provides UniAgents to manage the lifecycle of data collection plug-ins. For more information, see **Connecting to AOM** and **(New) Connecting to AOM**.

Dashboards

AOM provides more than 30 out-of-the-box dashboards and also enables you to customize your own dashboards. The dashboards support multiple graphs, such as tables, pies, maps, and line graphs. You can add custom variables and filters to screen metrics and logs for operations analysis. For more information, see **Dashboard Monitoring**.

Alarm Center

AOM centrally manages Prometheus and log alarm rules, alarm templates, and notification objects. It has more than 100 preset alarm templates and supports alarm grouping, suppression, and silence to reduce noise. You can customize message templates and receive notifications through emails, SMS, voice, WeCom, DingTalk, Lark, and Webhook. For more information, see **Alarm Monitoring**.

Log Management

AOM provides all-scenario log collection, search of tens of billions of logs in seconds, PB-level storage, one-stop log processing, visualized log SQL analysis, and log transfer capabilities, meeting application O&M, security compliance, and operations analysis requirements. For more information, see **(New) Log Management** and **(Old) Log Management**.

Prometheus Monitoring

AOM is interconnected with the open-source Prometheus ecosystem to provide fully hosted Prometheus services. It also supports PromQL query and metric browsing. Prometheus data can be remotely read and written. In addition, it supports interconnection with Grafana. For more information, see **Prometheus Monitoring**.

Infrastructure Monitoring

AOM collects monitoring data from hosts (processes) and containers (such as workloads and pods), displays infrastructure data (such as CPU/memory/disk usage and logs of hosts), and associates monitoring data for analysis. For more information, see **Infrastructure Monitoring**.

Application Performance & User Experience Monitoring

AOM analyzes the latency and exceptions from webs/apps to application microservices/databases/middleware, enabling O&M personnel to quickly locate root causes and identify application performance bottlenecks, thereby protecting user experience. For more information, see **Intelligent Insights**.

Settings

AOM provides service authorization, authentication management, global configuration, data subscription, and collection settings. Through service authorization, you can grant the current user the permissions to access cloud service data in one click. Through authentication management, you can create an access code and configure API call permissions for the current user. Through global configuration, you can enable or disable metric collection and Tag Management Service (TMS) tag display. Through data subscription, you can subscribe to metrics or alarms. Through collection settings, you can manage UniAgents, collection plug-ins, collectors in CCE clusters, host groups, and proxy areas, and check the operation logs of UniAgents and collection plug-ins. For more information, see **Global Settings** and **Managing Collector Base UniAgent**.

6 Comparison Between AOM 1.0 and AOM 2.0

Based on AOM 1.0 functions and common application monitoring, AOM 2.0 collects and monitors more metrics and log data, and displays monitoring results in a visualized manner.

This section compares AOM 1.0 with AOM 2.0.

Table 6-1 Comparison between AOM 1.0 and AOM 2.0

Function		Description	AOM 1.0	AOM 2.0
Resource monitoring	Access center	Quickly ingest metrics at the business, application, middleware, and infrastructure layers for monitoring.	Not supported.	Supported.
	Dashboard	Resource metrics, logs, and performance data are displayed in multiple graphs on the same screen.	Partially supported. Only metric data and system performance data can be monitored in a visualized manner.	Supported.
	Alarm management	You can set event conditions for services or set threshold criteria for resource metrics. When an alarm is generated due to an exception in AOM or a related service, the alarm information is sent to the specified personnel by email, SMS, or WeCom.	Partially supported. During alarm rule creation, metrics can be selected by metric type or running Prometheus commands, but cannot be selected from full metrics.	Supported.

Function		Description	AOM 1.0	AOM 2.0
	Application insights	Based on the CMDB model, the health status of resources is monitored by layer from applications, components, to environments.	Partially supported. The CMDB model is not supported. Metric data is monitored by layer.	Supported.
	Container insights	AOM monitors CCE resource usage, status, and alarms from workload and cluster dimensions for fast response and smooth workload running.	Supported.	Supported.
	Metric browsing	You can monitor metric data and trends of each resource and log data in real time, and create alarm rules for metrics to view services and analyze associated data in real time.	Partially supported. Only metric data can be monitored and analyzed.	Supported.
	Infrastructure monitoring	The running status of hosts and cloud services, and VM CPU, memory, and disk information can be monitored in real time.	Supported.	Supported.
	Prometheus monitoring	AOM is fully interconnected with the open-source Prometheus ecosystem, monitors various components, and flexibly expands cloud native component metric plug-ins.	Not supported.	Supported.
	Business monitoring	ELB log data reported to LTS are extracted as metrics for unified management. This facilitates real-time monitoring on the metric browsing and dashboard pages.	Not supported.	Supported.
	Log analysis	You can quickly search for required logs from massive quantities of logs. You can also quickly locate faults by analyzing the log source and context.	Supported.	Supported.

Function		Description	AOM 1.0	AOM 2.0
	Process monitoring	Rules can be set to discover deployed applications and collect associated metrics. Drill-down (from applications to components, instances, and containers) is also supported. Applications and components can be monitored from multiple dimensions.	Supported.	Supported.
	Data subscription	AOM allows you to subscribe to metrics or alarms. After subscription, data can be forwarded to custom Kafka or Distributed Message Service (DMS) topics for retrieval.	Supported.	Supported.
	Collection management	You can use UniAgents to schedule collection tasks to collect data.	Not supported.	Supported.

As functions of AOM 1.0 are gradually replaced by those of AOM 2.0, AOM 1.0 will be brought offline soon. You are advised to upgrade AOM 1.0 to AOM 2.0. For details, see **Upgrading to AOM 2.0**.

7 Relationships Between AOM and Other Services

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

SMN

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

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

OBS

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

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

LTS

Log Tank Service (LTS) can collect, analyze, and store log data. You can use LTS for efficient device O&M, service trend analysis, security audits, and monitoring.

AOM is a unified entry for Huawei Cloud observability analysis. It does not provide log functions, but integrates them from LTS.

CTS

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

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

IAM

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

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

Cloud Eye

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

AOM calls Cloud Eye APIs to obtain monitoring data of cloud services and displays them on the console so that you can monitor these services centrally.

APM

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

AOM incorporates APM functions for unified O&M. APM also has its own independent console and can be used separately.

VPC

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

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

ELB

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

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

RDS

RDS is a cloud-based web service which is reliable, scalable, easy to manage, and ready to use out-of-the-box.

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

DCS

DCS is an online, distributed, in-memory cache service compatible with Redis, Memcached, and In-Memory Data Grid (IMDG). It is reliable, scalable, ready to use out-of-the-box, and easy to manage, meeting your requirements for high read/write performance and fast data access.

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

CCE

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

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

CCI

CCI is a serverless container engine that allows you to run containers without creating and managing server clusters.

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

ServiceStage

ServiceStage is a one-stop PaaS service that provides cloud-based application hosting, simplifying application lifecycle management, from deployment, monitoring, O&M, to governance. It provides a microservice framework compatible with mainstream open-source ecosystems and enables quick building of distributed applications.

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

FunctionGraph

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

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

IEF

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

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

ECS

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

When purchasing an ECS, ensure that its OS meets the requirements in **Table 9-2**. In addition, install a UniAgent on the ECS. Otherwise, the ECS cannot be monitored by AOM. You can monitor basic resources, applications, logs, and alarms about this ECS on the AOM console.

BMS

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

When purchasing a BMS server, ensure that its OS meets the requirements in **Table 9-2**. In addition, install a UniAgent on the server. Otherwise, the server cannot be monitored by AOM. You can monitor basic resources, applications, logs, and alarms about this server on the AOM console.

8 Comparison Between AOM 2.0 and Cloud Eye

This section compares the cloud service monitoring functions of AOM 2.0 and Cloud Eye.

AOM metric data comes from Cloud Eye. AOM's metric data is in Prometheus format while Cloud Eye's metric data is in a custom format. **Table 8-1** compares the cloud service monitoring functions of AOM and Cloud Eye.

Table 8-1 Comparing the cloud service monitoring functions of AOM and Cloud Eye

Function	Cloud Eye	AOM 2.0
Unified monitoring across accounts	Not supported.	Supported.
Customization of data storage duration	Not supported (default: 3 months).	Supported (15 days, 30 days, 60 days, or 90 days).
Data export	• Aggregated data of the last three months can be exported. • Raw data of the last 48 hours can be exported.	• Dashboards and APIs can be exported. • Data can be written to Kafka.
Aggregate query	Only simple query is supported.	Multi-instance aggregation query is supported. For example, aggregation by tag or resource group.
PromQL syntax	Not supported.	Supported when you use alarm rules and dashboards, and browse metrics.

Function	Cloud Eye	AOM 2.0
Dashboards	Single-instance dashboards are supported for standard cloud products.	Various preset templates are provided.
Graph types supported by dashboards	2	8+
Monitoring views supported by a dashboard	50	100+
Alarm rules that can be created	Max.: 1000.	Default: 3000+. More than 10,000 rules can be supported.
Alarm rules that can be added to an alarm template	Max.: 50.	More than 20 cloud services can be added, and more than 100 alarm rules can be added for each cloud service.
Time that the alarm history can be kept	7 days.	1 year.
Objects that can be selected for single alarm rule creation	5000	Not limited. You can select all resources, and implement regular expression or exact match.
Alarm aggregation	Not supported.	Alarm aggregation based on PromQL syntax is supported. For example, implement alarm aggregation based on the total Content Delivery Network (CDN) bandwidth, total Elastic IP (EIP) bandwidth, or Object Storage Service (OBS) storage condition.
Connecting to the on-premises Grafana	Not supported.	Prometheus data sources can be directly connected to on-premises Grafana.

Function	Cloud Eye	AOM 2.0
Interconnecting with on-premises self-built Prometheus	Not supported.	Data can be directly written to self-built Prometheus.
Business monitoring	Not supported.	Business monitoring based on Prometheus, LTS logs, and custom channels is supported.
On-premises IDC monitoring	Not supported.	Prometheus Exporter-based on-premises hardware, storage, and network monitoring is supported.
On-premises middleware monitoring	Not supported.	On-premises middleware such as MongoDB, Redis, and RocketMQ can be monitored.

9 Restrictions

Resource Monitoring Restrictions

Table 9-1 Resource monitoring restrictions

Category	Item	Description
Dashboards	Dashboards	A maximum of 1,000 dashboards can be created in a region.
	Graphs in a dashboard	A maximum of 50 graphs can be added to a dashboard.
	Resources, threshold rules, components, or hosts in a graph	• A maximum of 12 resources can be added to a digit graph. Only one resource can be displayed. By default, the first resource is displayed. • A maximum of ten threshold rules can be added to a threshold status graph. • A maximum of ten hosts can be added to a host status graph. • A maximum of ten components can be added to a component status graph.
Metrics	Metric storage duration	• Metric data can be stored for up to 30 days. • ICAgent collects data at an interval of one minute. This interval cannot be changed.
	Storage duration of associated metric items	After resources (such as clusters, components, and hosts) are deleted, their metric items can still be stored for up to 30 days.
	Metric dimensions	A maximum of 20 dimensions can be configured for a metric.

Category	Item	Description
	Metrics queried using the 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 reported	A single request cannot exceed 40 KB. The timestamp of a reported metric cannot be 10 minutes later than the standard UTC time. In addition, out-of-order metrics are not received. That is, if a metric is reported at a certain time point, the metrics of earlier time points cannot be reported.
	Application metric	• When the number of containers on a host exceeds 1,000, 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 within 1,000, the ICAgent resumes the collection of application metrics and the **ICAgent Stopped Collecting Application Metrics** alarm is cleared.
	Job metric	A job automatically exits after it is completed. To monitor metrics of a job, ensure that its survival time is greater than 90s so that the ICAgent can collect its metric data.
	Resources consumed by the ICAgent	When the ICAgent collects basic metrics, the resources consumed by the ICAgent are related to 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 fewer than 1,000 containers run on a single node.

Category	Item	Description
	Metric dimension format	Metric dimension tags must comply with the AOM or Prometheus standard format so that the metrics can be reported to AOM. • AOM standard format: Only letters, digits, and underscores (_) are allowed. Start with a letter or underscore. • Prometheus standard format: Only ASCII letters, digits, and underscores (_) are allowed. The following regular expression must be met: [a-zA-Z_][a-zA-Z0-9_]*.
Alarm rules	Alarm rules	A maximum of 3,000 alarm rules (including metric alarm rules and event alarm rules) can be created.
	Alarm templates	A maximum of 150 alarm templates can be created.
Logs	Restrictions on the log function	For more information, see **LTS Usage Restrictions**.
	Log file	Only text log files can be collected. Other types of log files (for example, 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 1,000 standard container output log files. These files must be in JSON format.
	Resources consumed during log file collection	The resources consumed during log file collection are closely related to the log volume, number of files, network bandwidth, and backend service processing capability.
	Log discarding	When a single log line exceeds 10,240 bytes, it will be discarded.

Category	Item	Description
	Log collection path	**Linux** • Collection paths support recursion. You can use double asterisks (**) to collect logs from up to five directory levels. Example: **/var/logs/**/a.log** • Collection paths support fuzzy match. You can use an asterisk (*) to represent one or more characters of a directory or file name. Example: **/var/logs/*/a.log** or **/var/logs/service/a*.log** • If the collection path is set to a directory, for example, **/var/logs/**, only **.log**, **.trace**, and **.out** files in the directory are collected. If the collection path is set to name of a text file, that file is directly collected. • Each collection path must be unique. That is, the same path of the same host cannot be configured for different log groups and log streams. **Windows** • Collection paths support recursion. You can use double asterisks (**) to collect logs from up to five directory levels. Example: **C:\var\service**\a.log** • Collection paths support fuzzy match. You can use an asterisk (*) to represent one or more characters of a directory or file name. Examples: **C:\var\service*\a.log** and **C:\var\service\a*.log** • Each collection path must be unique. That is, the same path of the same host cannot be configured for different log groups and log streams.
	Log repetition	When the ICAgent is restarted, identical data may be collected around the restart time.
	Historical logs	The storage duration and prices of log data vary according to editions. For details, see **AOM Pricing Details**.
Alarm list	Time range for alarm query	You can query alarms generated within 31 days in the last year.
	Time range for event query	You can query events generated within 31 days in the last year.

Category	Item	Description
Application discovery	Application discovery rules	A maximum of 100 application discovery rules can be created.

Collection Management Restrictions

- OS Restrictions
  - For Linux x86_64 hosts, all the OSs and versions listed in **Table 9-2** are supported.
  - For Linux Arm hosts, CentOS 7.4/7.5/7.6, EulerOS 2.0, and Ubuntu 18.04 are supported.

Table 9-2 Linux OSs and versions supported by UniAgent

OS	Version				
Euler OS	1.1 64-bit	2.0 64-bit			
Cent OS	7.1 64-bit	7.2 64-bit	7.3 64-bit	7.4 64-bit	7.5 64-bit
	7.6 64-bit	7.7 64-bit	7.8 64-bit	7.9 64-bit	8.0 64-bit
Ubuntu	16.04 server 64-bit	18.04 server 64-bit	20.04 server 64-bit	22.04 server 64-bit	

Table 9-3 Windows OSs and versions supported by UniAgent

OS	Version
Windows Server	Windows Server 2012 R2 Standard 64-bit
	Windows Server 2012 R2 Standard English 64-bit
	Windows Server 2012 R2 Datacenter 64-bit
	Windows Server 2012 R2 Datacenter English 64-bit
	Windows Server 2016 Standard 64-bit
	Windows Server 2016 Standard English 64-bit
	Windows Server 2016 Datacenter 64-bit
	Windows Server 2016 Datacenter English 64-bit
	Windows Server 2019 Standard 64-bit

OS	Version
	Windows Server 2019 Standard English 64-bit
	Windows Server 2019 Datacenter 64-bit
	Windows Server 2019 Datacenter English 64-bit

- Resource Restrictions

Table 9-4 Resource restrictions

Item	Description
UniAgent client	When the average CPU usage is greater than 50% or the memory is greater than 100 MB for two minutes, the UniAgent client automatically restarts.
Installing, upgrading, or uninstalling UniAgents	You can install, upgrade, or uninstall UniAgents for a maximum of 100 hosts at a time.
Deleting hosts	You can delete a maximum of 50 hosts where UniAgents are not installed, offline, or fail to be installed at a time.

10 Metric Overview

10.1 Introduction

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

NOTE

This section describes only the metrics defined by the collection plug-in. Metrics reported by other cloud services or APIs are not included.

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 10-1**.

  Table 10-1 Namespaces of system metrics

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

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

Metric Dimensions

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

- Dimensions of system metrics are fixed. Different types of metrics have different dimensions. For details, see **10.12 Metric Dimensions**.
- Dimensions of custom metrics must be 1 to 32 characters long, which need to be customized.

10.2 Basic Metrics: VM Metrics

This section describes the categories, names, and meanings of VM metrics reported by ICAgents to AOM.

- If the host type is **CCE**, you can view disk partition metrics. The supported OSs are CentOS 7.6 and EulerOS 2.5.
- Log in to the CCE node as the **root** user and run the **docker info | grep 'Storage Driver'** command to check the Docker storage driver type. If the command output shows driver type **Device Mapper**, thin pool metrics can be viewed. Otherwise, thin pool metrics cannot be viewed.
- 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.
- For the total and used physical disk space, only the space of the local disk partitions' file systems is counted. The file systems (such as JuiceFS, NFS, and SMB) mounted to the host through the network are not taken into account.
- Cluster metrics are aggregated by AOM based on host metrics, and do not include the metrics of master hosts.

Table 10-2 VM metrics

Category	Metric	Metric Name	Description	Value Range	Unit
Network metrics	aom_node_network_receive_bytes	Downlink Rate (BPS)	Inbound traffic rate of a measured object	≥ 0	Bytes/s
	aom_node_network_receive_packets	Downlink Rate (PPS)	Number of data packets received by a NIC per second	≥ 0	Packets/s

Category	Metric	Metric Name	Description	Value Range	Unit
	aom_node_network_receive_error_packets	Downlink Error Rate	Number of error packets received by a NIC per second	≥ 0	Packets/s
	aom_node_network_transmit_bytes	Uplink Rate (BPS)	Outbound traffic rate of a measured object	≥ 0	Bytes/s
	aom_node_network_transmit_error_packets	Uplink Error Rate	Number of error packets transmitted by a NIC per second	≥ 0	Packets/s
	aom_node_network_transmit_packets	Uplink Rate (PPS)	Number of data packets transmitted by a NIC per second	≥ 0	Packets/s
	aom_node_network_total_bytes	Total Rate (BPS)	Total inbound and outbound traffic rate of a measured object	≥ 0	Bytes/s
Disk metrics	aom_node_disk_read_kilobytes	Disk Read Rate	Volume of data read from a disk per second	≥ 0	KB/s
	aom_node_disk_write_kilobytes	Disk Write Rate	Volume of data written into a disk per second	≥ 0	KB/s
Disk partition metrics	aom_host_diskpartition_thinpool_metadata_percent	Thin Pool Metadata Space Usage	Percentage of the thin pool's used metadata space to the total metadata space on a CCE node	0–100	%
	aom_host_diskpartition_thinpool_data_percent	Thin Pool Data Space Usage	Percentage of the thin pool's used data space to the total data space on a CCE node	0–100	%
	aom_host_diskpartition_total_capacity_megabytes	Thin Pool Disk Partition Space	Total thin pool disk partition space on a CCE node	≥ 0	MB

Category	Metric	Metric Name	Description	Value Range	Unit
File system metrics	aom_node_disk_available_capacity_megabytes	Available Disk Space	Disk space that has not been used	≥ 0	MB
	aom_node_disk_capacity_megabytes	Total Disk Space	Total disk space	≥ 0	MB
	aom_node_disk_rw_status	Disk Read/ Write Status	Read or write status of a disk	0 or 1 • 0: read/ write • 1: read-only	N/A
	aom_node_disk_usage	Disk Usage	Percentage of the used disk space to the total disk space	0–100	%
Host metrics	aom_node_cpu_limit_core	Total CPU Cores	Total number of CPU cores that have been applied for a measured object	≥ 1	Cores
	aom_node_cpu_used_core	Used CPU Cores	Number of CPU cores used by a measured object	≥ 0	Cores
	aom_node_cpu_usage	CPU Usage	CPU usage of a measured object	0–100	%
	aom_node_memory_free_megabytes	Available Physical Memory	Available physical memory of a measured object	≥ 0	MB
	aom_node_virtual_memory_free_megabytes	Available Virtual Memory	Available virtual memory of a measured object	≥ 0	MB
	aom_node_gpu_memory_free_megabytes	GPU Memory	GPU memory of a measured object	> 0	MB
	aom_node_gpu_memory_usage	GPU Memory Usage	Percentage of the used GPU memory to the total GPU memory	0–100	%

Category	Metric	Metric Name	Description	Value Range	Unit
	aom_node_gpu_memory_used_megabytes	Used GPU Memory	GPU memory used by a measured object	≥ 0	MB
	aom_node_gpu_usage	GPU Usage	GPU usage of a measured object	0–100	%
	aom_node_npu_memory_free_megabytes	NPU Memory	NPU memory of a measured object **Only NPU metrics of CCE hosts can be collected.**	> 0	MB
	aom_node_npu_memory_usage	NPU Memory Usage	Percentage of the used NPU memory to the total NPU memory **Only NPU metrics of CCE hosts can be collected.**	0–100	%
	aom_node_npu_memory_used_megabytes	Used NPU Memory	NPU memory used by a measured object **Only NPU metrics of CCE hosts can be collected.**	≥ 0	MB
	aom_node_npu_usage	NPU Usage	NPU usage of a measured object **Only NPU metrics of CCE hosts can be collected.**	0–100	%
	aom_node_npu_temperature_centigrade	NPU Temperature	NPU temperature of a measured object **Only NPU metrics of CCE hosts can be collected.**	-	°C

Category	Metric	Metric Name	Description	Value Range	Unit
	aom_node_memory_usage	Physical Memory Usage	Percentage of the used physical memory to the total physical memory applied for a measured object	0–100	%
	aom_node_ntp_offset_ms	NTP Offset	Offset between the local time of the host and the NTP server time. The closer the NTP offset is to 0, the closer the local time of the host is to the time of the NTP server.	-	ms
	aom_node_ntp_server_status	NTP Server Status	Whether the host is connected to the NTP server	0 or 1 • 0: Connected • 1: Not connected	N/A
	aom_node_ntp_status	NTP Synchronization Status	Whether the local time of the host is synchronized with the NTP server time	0 or 1 • 0: Synchronous • 1: Asynchronous	N/A
	aom_node_process_number	Processes	Number of processes on a measured object	≥ 0	Count
	aom_node_gpu_temperature_centigrade	GPU Temperature	GPU temperature of a measured object	-	°C

Category	Metric	Metric Name	Description	Value Range	Unit
	aom_node_memory_total_megabytes	Total Physical Memory	Total physical memory that has been applied for a measured object	≥ 0	MB
	aom_node_virtual_memory_total_megabytes	Total Virtual Memory	Total virtual memory that has been applied for a measured object	≥ 0	MB
	aom_node_virtual_memory_usage	Virtual Memory Usage	Percentage of the used virtual memory to the total virtual memory	0–100	%
	aom_node_current_threads_num	Current Threads	Number of threads created on a host	≥ 0	Count
	aom_node_sys_max_threads_num	Max. Threads	Maximum number of threads that can be created on a host	≥ 0	Count
	aom_node_phy_disk_total_capacity_megabytes	Total Physical Disk Space	Total disk space of a host	≥ 0	MB
	aom_node_physical_disk_total_used_megabytes	Used Physical Disk Space	Used disk space of a host	≥ 0	MB
	aom_billing_hostUsed	Hosts	Number of hosts connected per day	≥ 0	Count
Cluster metrics	aom_cluster_cpu_limit_core	Total CPU Cores	Total number of CPU cores that have been applied for a measured object	≥ 1	Cores
	aom_cluster_cpu_used_core	Used CPU Cores	Number of CPU cores used by a measured object	≥ 0	Cores

Category	Metric	Metric Name	Description	Value Range	Unit
	aom_cluster_cpu_usage	CPU Usage	CPU usage of a measured object	0–100	%
	aom_cluster_disk_available_capacity_megabytes	Available Disk Space	Disk space that has not been used	≥ 0	MB
	aom_cluster_disk_capacity_megabytes	Total Disk Space	Total disk space	≥ 0	MB
	aom_cluster_disk_usage	Disk Usage	Percentage of the used disk space to the total disk space	0–100	%
	aom_cluster_memory_free_megabytes	Available Physical Memory	Available physical memory of a measured object	≥ 0	MB
	aom_cluster_virtual_memory_free_megabytes	Available Virtual Memory	Available virtual memory of a measured object	≥ 0	MB
	aom_cluster_gpu_memory_free_megabytes	Available GPU Memory	Available GPU memory of a measured object	> 0	MB
	aom_cluster_gpu_memory_usage	GPU Memory Usage	Percentage of the used GPU memory to the total GPU memory	0–100	%
	aom_cluster_gpu_memory_used_megabytes	Used GPU Memory	GPU memory used by a measured object	≥ 0	MB
	aom_cluster_gpu_usage	GPU Usage	GPU usage of a measured object	0–100	%
	aom_cluster_memory_usage	Physical Memory Usage	Percentage of the used physical memory to the total physical memory applied for a measured object	0–100	%

Category	Metric	Metric Name	Description	Value Range	Unit
	aom_cluster_network_receive_bytes	Downlink Rate (BPS)	Inbound traffic rate of a measured object	≥ 0	Bytes/s
	aom_cluster_network_transmit_bytes	Uplink Rate (BPS)	Outbound traffic rate of a measured object	≥ 0	Bytes/s
	aom_cluster_memory_total_megabytes	Total Physical Memory	Total physical memory that has been applied for a measured object	≥ 0	MB
	aom_cluster_virtual_memory_total_megabytes	Total Virtual Memory	Total virtual memory of a measured object	≥ 0	MB
	aom_cluster_virtual_memory_usage	Virtual Memory Usage	Percentage of the used virtual memory to the total virtual memory	0–100	%
Container metrics	aom_container_cpu_limit_core	Total CPU Cores	Total number of CPU cores restricted for a measured object	≥ 1	Cores
	aom_container_cpu_used_core	Used CPU Cores	Number of CPU cores used by a measured object	≥ 0	Cores
	aom_container_cpu_usage	CPU Usage	CPU usage of a measured object Percentage of the used CPU cores to the total CPU cores restricted for a measured object	0–100	%
	aom_container_disk_read_kilobytes	Disk Read Rate	Volume of data read from a disk per second	≥ 0	KB/s
	aom_container_disk_write_kilobytes	Disk Write Rate	Volume of data written into a disk per second	≥ 0	KB/s

Category	Metric	Metric Name	Description	Value Range	Unit
	aom_container_filesystem_available_capacity_megabytes	Available File System Capacity	Available file system capacity of a measured object. This metric is available only for containers using the Device Mapper storage drive in the Kubernetes cluster of version 1.11 or later.	≥ 0	MB
	aom_container_filesystem_capacity_megabytes	Total File System Capacity	Total file system capacity of a measured object. This metric is available only for containers using the Device Mapper storage drive in the Kubernetes cluster of version 1.11 or later.	≥ 0	MB
	aom_container_filesystem_usage	File System Usage	File system usage of a measured object. That is, the percentage of the used file system to the total file system. This metric is available only for containers using the Device Mapper storage drive in the Kubernetes cluster of version 1.11 or later.	0–100	%
	aom_container_gpu_memory_free_megabytes	GPU Memory	GPU memory of a measured object	> 0	MB

Category	Metric	Metric Name	Description	Value Range	Unit
	aom_container_gpu_memory_usage	GPU Memory Usage	Percentage of the used GPU memory to the total GPU memory	0–100	%
	aom_container_gpu_memory_used_megabytes	Used GPU Memory	GPU memory used by a measured object	≥ 0	MB
	aom_container_gpu_usage	GPU Usage	GPU usage of a measured object	0–100	%
	aom_container_npu_memory_free_megabytes	NPU Memory	NPU memory of a measured object	> 0	MB
	aom_container_npu_memory_usage	NPU Memory Usage	Percentage of the used NPU memory to the total NPU memory	0–100	%
	aom_container_npu_memory_used_megabytes	Used NPU Memory	NPU memory used by a measured object	≥ 0	MB
	aom_container_npu_usage	NPU Usage	NPU usage of a measured object	0–100	%
	aom_container_memory_request_megabytes	Total Physical Memory	Total physical memory restricted for a measured object	≥ 0	MB
	aom_container_memory_usage	Physical Memory Usage	Percentage of the used physical memory to the total physical memory restricted for a measured object	0–100	%
	aom_container_memory_used_megabytes	Used Physical Memory	Physical memory (resident set size) used by a measured object	≥ 0	MB

Category	Metric	Metric Name	Description	Value Range	Unit
	aom_container_network_receive_bytes	Downlink Rate (BPS)	Inbound traffic rate of a measured object	≥ 0	Bytes/s
	aom_container_network_receive_packets	Downlink Rate (PPS)	Number of data packets received by a NIC per second	≥ 0	Packets/s
	aom_container_network_receive_error_packets	Downlink Error Rate	Number of error packets received by a NIC per second	≥ 0	Packets/s
	aom_container_network_rx_error_packets	Error Packets Received	Number of error packets received by a measured object	≥ 0	Count
	aom_container_network_transmit_bytes	Uplink Rate (BPS)	Outbound traffic rate of a measured object	≥ 0	Bytes/s
	aom_container_network_transmit_error_packets	Uplink Error Rate	Number of error packets transmitted by a NIC per second	≥ 0	Packets/s
	aom_container_network_transmit_packets	Uplink Rate (PPS)	Number of data packets transmitted by a NIC per second	≥ 0	Packets/s
	aom_container_memory_workingset_usage	Working Set Memory Usage	Usage of the working set memory	0–100	%
	aom_container_memory_workingset_used_megabytes	Used Working Set Memory	Working set memory that has been used	≥ 0	MB
Process metrics	aom_process_cpu_limit_core	Total CPU Cores	Total number of CPU cores that have been applied for a measured object	≥ 1	Cores
	aom_process_cpu_used_core	Used CPU Cores	Number of CPU cores used by a measured object	≥ 0	Cores

Category	Metric	Metric Name	Description	Value Range	Unit
	aom_process_cpu_usage	CPU Usage	CPU usage of a measured object Percentage of the used CPU cores to the CPU cores that have been applied	0–100	%
	aom_process_handle_count	Handles	Number of handles used by a measured object	≥ 0	Count
	aom_process_max_handle_count	Max. Handles	Maximum number of handles used by a measured object	≥ 0	Count
	aom_process_memory_request_megabytes	Total Physical Memory	Total physical memory that has been applied for a measured object	≥ 0	MB
	aom_process_memory_usage	Physical Memory Usage	Percentage of the used physical memory to the total physical memory applied for a measured object	0–100	%
	aom_process_memory_used_megabytes	Used Physical Memory	Physical memory (resident set size) used by a measured object	≥ 0	MB
	aom_process_thread_count	Threads	Number of threads used by a measured object	≥ 0	Count
	aom_process_virtual_memory_total_megabytes	Total Virtual Memory	Total virtual memory that has been applied for a measured object	≥ 0	MB

10.3 Basic Metrics: Container Metrics

This section describes the categories, names, and meanings of metrics reported to AOM from CCE's kube-prometheus-stack add-on or on-premises Kubernetes clusters.

Table 10-3 Metrics of containers running in CCE or on-premises Kubernetes clusters

Target Name	Job Name	Metric	Description
• serviceMonitor/monitoring/coredns/0 • serviceMonitor/monitoring/node-local-dns/0	coredns and node-local-dns	coredns_build_info	Information to build CoreDNS
		coredns_cache_entries	Number of entries in the CoreDNS cache
		coredns_cache_size	CoreDNS cache size
		coredns_cache_hits_total	Number of CoreDNS cache hits
		coredns_cache_misses_total	Number of CoreDNS cache misses
		coredns_cache_requests_total	Total number of CoreDNS resolution requests in different dimensions
		coredns_dns_request_duration_seconds_bucket	CoreDNS request latency
		coredns_dns_request_duration_seconds_count	CoreDNS request processing time (seconds)
		coredns_dns_request_duration_seconds_sum	Total CoreDNS request processing time (seconds)
		coredns_dns_request_size_bytes_bucket	Size of the CoreDNS request in bytes
		coredns_dns_request_size_bytes_count	CoreDNS request byte count
		coredns_dns_request_size_bytes_sum	Total CoreDNS request bytes
		coredns_dns_requests_total	Total number of CoreDNS requests

Target Name	Job Name	Metric	Description
		coredns_dns_response_size_bytes_bucket	Size of the returned CoreDNS response in bytes
		coredns_dns_response_size_bytes_count	CoreDNS response byte count
		coredns_dns_response_size_bytes_sum	Total CoreDNS response bytes
		coredns_dns_responses_total	Total number of CoreDNS response codes
		coredns_forward_conn_cache_hits_total	Total number of cache hits for each protocol and data flow
		coredns_forward_conn_cache_misses_total	Total number of cache misses for each protocol and data flow
		coredns_forward_healthcheck_broken_total	Total forwarding health check failures
		coredns_forward_healthcheck_failures_total	Total forwarding health check faults
		coredns_forward_max_concurrent_rejects_total	Total number of requests rejected due to excessive concurrent requests
		coredns_forward_request_duration_seconds_bucket	CoreDNS forwarding request latency
		coredns_forward_request_duration_seconds_count	CoreDNS forwarding request duration in seconds
		coredns_forward_request_duration_seconds_sum	Total CoreDNS forwarding request duration in seconds
		coredns_forward_requests_total	Total number of requests for each data flow
		coredns_forward_responses_total	Total number of responses to each data flow

Target Name	Job Name	Metric	Description
		coredns_health_request_duration_seconds_bucket	CoreDNS health check request latency
		coredns_health_request_duration_seconds_count	CoreDNS health check request duration in seconds
		coredns_health_request_duration_seconds_sum	Total CoreDNS health check request duration in seconds
		coredns_health_request_failures_total	Total number of failed CoreDNS health check requests
		coredns_hosts_reload_timestamp_seconds	Timestamp of CoreDNS's last reload of the host file
		coredns_kubernetes_dns_programming_duration_seconds_bucket	DNS programming latency
		coredns_kubernetes_dns_programming_duration_seconds_count	DNS programming duration in seconds
		coredns_kubernetes_dns_programming_duration_seconds_sum	Total DNS programming duration in seconds
		coredns_local_localhost_requests_total	Total number of localhost requests processed by CoreDNS
		coredns_nodecache_setup_errors_total	Total number of node cache plug-in setting errors
		coredns_dns_response_rcode_count_total	Cumulative count of response codes
		coredns_dns_request_count_total	Cumulative count of DNS requests made per zone, protocol, and family
		coredns_dns_request_do_count_total	Cumulative count of requests with the DO bit set
		coredns_dns_do_requests_total	Number of requests with the DO bit set

Target Name	Job Name	Metric	Description
		coredns_dns_request_type_count_total	Cumulative count of DNS requests per type
		coredns_panics_total	Total number of CoreDNS abnormal exits
		coredns_plugin_enabled	Whether a plugin is enabled in CoreDNS
		coredns_reload_failed_total	Total number of configuration files that fail to be reloaded
serviceMonitor/monitoring/kube-apiserver/0	apiserver	aggregator_unavailable_apiservice	Number of unavailable APIServices
		apiserver_admission_controller_admission_duration_seconds_bucket	Processing delay of an admission controller
		apiserver_admission_webhook_admission_duration_seconds_bucket	Processing delay of an admission webhook
		apiserver_admission_webhook_admission_duration_seconds_count	Number of admission webhook processing requests
		apiserver_client_certificate_expiration_seconds_bucket	Remaining validity period of the client certificate
		apiserver_client_certificate_expiration_seconds_count	Remaining validity period of the client certificate
		apiserver_current_inflight_requests	Number of read requests in process
		apiserver_request_duration_seconds_bucket	Delay of the client's access to the APIServer
		apiserver_request_total	Counter of API server requests broken out for code and other items

Target Name	Job Name	Metric	Description
		go_goroutines	Number of goroutines that exist
		kubernetes_build_info	Information to build Kubernetes
		process_cpu_seconds_total	Total process CPU time
		process_resident_memory_bytes	Size of the resident memory set
		rest_client_requests_total	Total number of HTTP requests, partitioned by status code and method
		workqueue_adds_total	Total number of additions handled by a work queue
		workqueue_depth	Current depth of a work queue
		workqueue_queue_duration_seconds_bucket	Duration that a task stays in the current queue
		aggregator_unavailable_apiservice_total	Number of unavailable APIServices
		rest_client_request_duration_seconds_bucket	Number of HTTP requests, partitioned by status code and method
serviceMonitor/monitoring/kubelet/0	kubelet	kubelet_certificate_manager_client_expiration_renew_errors	Number of certificate renewal errors
		kubelet_certificate_manager_client_ttl_seconds	Time-to-live (TTL) of the Kubelet client certificate
		kubelet_cgroup_manager_duration_seconds_bucket	Duration for destruction and update operations
		kubelet_cgroup_manager_duration_seconds_count	Number of destruction and update operations

Target Name	Job Name	Metric	Description
		kubelet_node_config_error	If a configuration-related error occurs on a node, the value of this metric is **true (1)**. If there is no configuration-related error, the value is **false (0)**.
		kubelet_node_name	Node name. The value is always **1**.
		kubelet_pleg_relist_duration_seconds_bucket	Duration for relisting pods in PLEG
		kubelet_pleg_relist_duration_seconds_count	Duration in seconds for relisting pods in PLEG
		kubelet_pleg_relist_interval_seconds_bucket	Interval between relisting operations in PLEG
		kubelet_pod_start_duration_seconds_count	Number of pods that have been started
		kubelet_pod_start_duration_seconds_bucket	Duration from the kubelet seeing a pod for the first time to the pod starting to run
		kubelet_pod_worker_duration_seconds_bucket	Duration for synchronizing a single pod.
		kubelet_running_containers	Number of running containers
		kubelet_running_pods	Number of running pods
		kubelet_runtime_operations_duration_seconds_bucket	Time of every operation
		kubelet_runtime_operations_errors_total	Number of errors in operations at runtime level
		kubelet_runtime_operations_total	Number of runtime operations of each type

Target Name	Job Name	Metric	Description
		kubelet_volume_stats_available_bytes	Number of available bytes in a volume
		kubelet_volume_stats_capacity_bytes	Capacity in bytes of a volume
		kubelet_volume_stats_inodes	Maximum number of inodes in a volume
		kubelet_volume_stats_inodes_used	Number of used inodes in a volume
		kubelet_volume_stats_used_bytes	Number of used bytes in a volume
		storage_operation_duration_seconds_bucket	Duration for each storage operation
		storage_operation_duration_seconds_count	Number of storage operations
		storage_operation_errors_total	Number of storage operation errors
		volume_manager_total_volumes	Number of volumes in Volume Manager
		rest_client_requests_total	Total number of HTTP requests, partitioned by status code and method
		rest_client_request_duration_seconds_bucket	Number of HTTP requests, partitioned by status code and method
		process_resident_memory_bytes	Size of the resident memory set
		process_cpu_seconds_total	Total process CPU time
		go_goroutines	Number of goroutines that exist
serviceMonitor/monitoring/kubelet/1	kubelet	container_cpu_cfs_periods_total	Total number of elapsed enforcement periods
		container_cpu_cfs_throttled_periods_total	Number of throttled periods

Target Name	Job Name	Metric	Description
		container_cpu_cfs_throttled_seconds_total	Total duration a container has been throttled
		container_cpu_load_average_10s	Value of container CPU load average over the last 10 seconds
		container_cpu_usage_seconds_total	Total CPU time consumed
		container_file_descriptors	Number of open file descriptors for a container
		container_fs_inodes_free	Number of available inodes in a file system
		container_fs_inodes_total	Total number of inodes in a file system
		container_fs_io_time_seconds_total	Cumulative time spent on doing I/Os by the disk or file system
		container_fs_limit_bytes	Total disk or file system capacity that can be consumed by a container
		container_fs_read_seconds_total	Total time a container spent on reading disk or file system data
		container_fs_reads_bytes_total	Cumulative amount of disk or file system data read by a container
		container_fs_reads_total	Cumulative number of disk or file system reads completed by a container
		container_fs_usage_bytes	File system usage
		container_fs_write_seconds_total	Total time a container spent on writing data to the disk or file system

Target Name	Job Name	Metric	Description
		container_fs_writes_bytes_total	Total amount of data written by a container to a disk or file system
		container_fs_writes_total	Cumulative number of disk or file system writes completed by a container
		container_memory_cache	Memory used for the page cache of a container
		container_memory_failcnt	Number of memory usage hits limits
		container_memory_max_usage_bytes	Maximum memory usage recorded for a container
		container_memory_rss	Size of the resident memory set for a container
		container_memory_swap	Container swap memory usage
		container_memory_usage_bytes	Current memory usage of a container
		container_memory_working_set_bytes	Memory usage of the working set of a container
		container_network_receive_bytes_total	Total volume of data received by a container network
		container_network_receive_errors_total	Cumulative number of errors encountered during reception
		container_network_receive_packets_dropped_total	Cumulative number of packets dropped during reception
		container_network_receive_packets_total	Cumulative number of packets received
		container_network_transmit_bytes_total	Total volume of data transmitted on a container network

Target Name	Job Name	Metric	Description
		container_network_transmit_errors_total	Cumulative number of errors encountered during transmission
		container_network_transmit_packets_dropped_total	Cumulative number of packets dropped during transmission
		container_network_transmit_packets_total	Cumulative number of packets transmitted
		container_spec_cpu_quota	CPU quota of a container
		container_spec_memory_limit_bytes	Memory limit for a container
		machine_cpu_cores	Number of CPU cores of the physical machine or VM
		machine_memory_bytes	Total memory size of the physical machine or VM
serviceMonitor/monitoring/kube-state-metrics/0	kube-state-metrics-prom	kube_cronjob_status_active	Whether the cronjob is actively running jobs
		kube_cronjob_info	Cronjob information
		kube_cronjob_labels	Label of a cronjob
		kube_configmap_info	ConfigMap information
		kube_daemonset_created	DaemonSet creation time
		kube_daemonset_status_current_number_scheduled	Number of DaemonSets that are being scheduled
		kube_daemonset_status_desired_number_scheduled	Number of DaemonSets expected to be scheduled
		kube_daemonset_status_number_available	Number of nodes that should be running a DaemonSet pod and have at least one DaemonSet pod running and available

Target Name	Job Name	Metric	Description
		kube_daemonset_status_number_misscheduled	Number of nodes that are not expected to run a DaemonSet pod
		kube_daemonset_status_number_ready	Number of nodes that should be running the DaemonSet pods and have one or more DaemonSet pods running and ready
		kube_daemonset_status_number_unavailable	Number of nodes that should be running the DaemonSet pods but have none of the DaemonSet pods running and available
		kube_daemonset_status_updated_number_scheduled	Number of nodes that are running an updated DaemonSet pod
		kube_deployment_created	Deployment creation timestamp
		kube_deployment_labels	Deployment labels
		kube_deployment_metadata_generation	Sequence number representing a specific generation of the desired state for a Deployment
		kube_deployment_spec_replicas	Number of desired replicas for a Deployment
		kube_deployment_spec_strategy_rollingupdate_max_unavailable	Maximum number of unavailable replicas during a rolling update of a Deployment
		kube_deployment_status_observed_generation	The generation observed by the Deployment controller
		kube_deployment_status_replicas	Number of current replicas of a Deployment

Target Name	Job Name	Metric	Description
		kube_deployment_status_replicas_available	Number of available replicas per Deployment
		kube_deployment_status_replicas_ready	Number of ready replicas per Deployment
		kube_deployment_status_replicas_unavailable	Number of unavailable replicas per Deployment
		kube_deployment_status_replicas_updated	Number of updated replicas per Deployment
		kube_job_info	Job information
		kube_namespace_labels	Namespace labels
		kube_node_labels	Node labels
		kube_node_info	Node information
		kube_node_spec_taint	Taint of a node
		kube_node_spec_unschedulable	Whether new pods can be scheduled to a node
		kube_node_status_allocatable	Allocatable resources on a node
		kube_node_status_capacity	Capacity for different resources on a node
		kube_node_status_condition	Node status condition
		kube_node_volcano_oversubscription_status	Node oversubscription status
		kube_persistentvolume_status_phase	PV status
		kube_persistentvolumeclaim_status_phase	PVC status
		kube_persistentvolume_info	PV information
		kube_persistentvolumeclaim_info	PVC information

Target Name	Job Name	Metric	Description
		kube_pod_container_info	Information about a container running in the pod
		kube_pod_container_resource_limits	Container resource limits
		kube_pod_container_resource_requests	Number of resources requested by a container
		kube_pod_container_status_last_terminated_reason	The last reason a container was in terminated state
		kube_pod_container_status_ready	Whether a container is in ready state
		kube_pod_container_status_restarts_total	Number of container restarts
		kube_pod_container_status_running	Whether a container is in running state
		kube_pod_container_status_terminated	Whether a container is in terminated state
		kube_pod_container_status_terminated_reason	The reason a container is in terminated state
		kube_pod_container_status_waiting	Whether a container is in waiting state
		kube_pod_container_status_waiting_reason	The reason a container is in waiting state
		kube_pod_info	Pod information
		kube_pod_labels	Pod labels
		kube_pod_owner	Object to which the pod belongs
		kube_pod_status_phase	Phase of the pod
		kube_pod_status_ready	Whether the pod is in ready state
		kube_secret_info	Secret information
		kube_statefulset_created	StatefulSet creation timestamp

Target Name	Job Name	Metric	Description
		kube_statefulset_labels	Information about StatefulSet labels
		kube_statefulset_metadata_generation	Sequence number representing a specific generation of the desired state for a StatefulSet
		kube_statefulset_replicas	Number of desired pods for a StatefulSet
		kube_statefulset_status_observed_generation	Generation observed by the StatefulSet controller
		kube_statefulset_status_replicas	Number of stateful replicas in a StatefulSet
		kube_statefulset_status_replicas_ready	Number of ready replicas in a StatefulSet
		kube_statefulset_status_replicas_updated	Number of updated replicas in a StatefulSet
		kube_job_spec_completions	Desired number of successfully finished pods that should run with the job
		kube_job_status_failed	Failed jobs
		kube_job_status_succeeded	Successful jobs
		kube_node_status_allocatable_cpu_cores	Number of allocatable CPU cores of a node
		kube_node_status_allocatable_memory_bytes	Total allocatable memory of a node
		kube_replicaset_owner	ReplicaSet owner.
		kube_resourcequota	Resource quota
		kube_pod_spec_volumes_persistentvolume-claims_info	Information about the PVC associated with the pod

Target Name	Job Name	Metric	Description
serviceMonitor/monitoring/prometheus-lightweight/0	prometheus-lightweight	vm_persistentqueue_blocks_dropped_total	Total number of dropped blocks in a send queue
		vm_persistentqueue_blocks_read_total	Total number of blocks read by a send queue
		vm_persistentqueue_blocks_written_total	Total number of blocks written to a send queue
		vm_persistentqueue_bytes_pending	Number of pending bytes in a send queue
		vm_persistentqueue_bytes_read_total	Total number of bytes read by a send queue
		vm_persistentqueue_bytes_written_total	Total number of bytes written to a send queue
		vm_promscrape_active_scrapers	Number of collected shards
		vm_promscrape_conn_read_errors_total	Total number of read errors during scrapes
		vm_promscrape_conn_write_errors_total	Total number of write errors during scrapes
		vm_promscrape_max_scrape_size_exceeded_errors_total	Total number of scrapes failed because responses exceed the size limit
		vm_promscrape_scrape_duration_seconds_sum	Time required for the scrape
		vm_promscrape_scrape_duration_seconds_count	Total time required for the scrape
		vm_promscrape_scrapes_total	Number of scrapes
		vmagent_remotewrite_bytes_sent_total	Total number of bytes sent through remote write
		vmagent_remotewrite_duration_seconds_sum	Time consumed by remote writes

Target Name	Job Name	Metric	Description
		vmagent_remotewrite_duration_seconds_count	Total time consumed by remote writes
		vmagent_remotewrite_packets_dropped_total	Total number of dropped packets during remote write
		vmagent_remotewrite_pending_data_bytes	Number of pending bytes during remote write
		vmagent_remotewrite_requests_total	Total number of remote write requests
		vmagent_remotewrite_retries_count_total	Total number of remote write retries
		go_goroutines	Number of goroutines that exist
serviceMonitor/monitoring/node-exporter/0	node-exporter	node_boot_time_seconds	Node boot time
		node_context_switches_total	Number of context switches
		node_cpu_seconds_total	Seconds the CPUs spent in each mode
		node_disk_io_now	Number of I/Os in progress
		node_disk_io_time_seconds_total	Total seconds spent doing I/Os
		node_disk_io_time_weighted_seconds_total	The weighted time spent doing I/Os
		node_disk_read_bytes_total	Number of bytes that are read
		node_disk_read_time_seconds_total	Number of seconds spent by all reads
		node_disk_reads_completed_total	Number of reads completed
		node_disk_write_time_seconds_total	Number of seconds spent by all writes
		node_disk_writes_completed_total	Number of writes completed
		node_disk_written_bytes_total	Number of bytes that are written

Target Name	Job Name	Metric	Description
		node_docker_thinpool_data_space_available	Available data space of a Docker thin pool
		node_docker_thinpool_metadata_space_available	Available metadata space of a Docker thin pool
		node_exporter_build_info	Node Exporter build information
		node_filefd_allocated	Allocated file descriptors
		node_filefd_maximum	Maximum number of file descriptors
		node_filesystem_avail_bytes	File system space that is available for use
		node_filesystem_device_error	Error in the mounted file system device
		node_filesystem_free_bytes	Remaining space of a file system
		node_filesystem_readonly	Read-only file system
		node_filesystem_size_bytes	Consumed space of a file system
		node_forks_total	Number of forks
		node_intr_total	Number of interruptions that occurred
		node_load1	1-minute average CPU load
		node_load15	15-minute average CPU load
		node_load5	5-minute average CPU load
		node_memory_Buffers_bytes	Memory of the node buffer
		node_memory_Cached_bytes	Memory for the node page cache
		node_memory_MemAvailable_bytes	Available memory of a node

Target Name	Job Name	Metric	Description
		node_memory_MemFree_bytes	Free memory of a node
		node_memory_MemTotal_bytes	Total memory of a node
		node_network_receive_bytes_total	Total amount of received data
		node_network_receive_drop_total	Total number of packets dropped during reception
		node_network_receive_errs_total	Total number of errors encountered during reception
		node_network_receive_packets_total	Total number of packets received
		node_network_transmit_bytes_total	Total number of sent bytes
		node_network_transmit_drop_total	Total number of dropped packets
		node_network_transmit_errs_total	Total number of errors encountered during transmission
		node_network_transmit_packets_total	Total number of packets sent
		node_procs_blocked	Blocked processes
		node_procs_running	Running processes
		node_sockstat_sockets_used	Number of sockets in use
		node_sockstat_TCP_alloc	Number of allocated TCP sockets
		node_sockstat_TCP_inuse	Number of TCP sockets in use
		node_sockstat_TCP_orphan	Number of orphaned TCP sockets
		node_sockstat_TCP_tw	Number of TCP sockets in the **TIME_WAIT** state
		node_sockstat_UDPLITE_inuse	Number of UDP-Lite sockets in use

Target Name	Job Name	Metric	Description
		node_sockstat_UDP_inuse	Number of UDP sockets in use
		node_sockstat_UDP_mem	UDP socket buffer usage
		node_timex_offset_seconds	Time offset
		node_timex_sync_status	Synchronization status of node clocks
		node_uname_info	System kernel information
		node_vmstat_oom_kill	Number of processes terminated due to insufficient memory
		process_cpu_seconds_total	Total process CPU time
		process_max_fds	Maximum number of file descriptors of a process
		process_open_fds	Opened file descriptors by a process
		process_resident_memory_bytes	Size of the resident memory set
		process_start_time_seconds	Process start time
		process_virtual_memory_bytes	Virtual memory size
		process_virtual_memory_max_bytes	Maximum available virtual memory capacity
		node_netstat_Tcp_ActiveOpens	Number of TCP connections that directly change from the **CLOSED** state to the **SYN-SENT** state
		node_netstat_Tcp_PassiveOpens	Number of TCP connections that directly change from the **LISTEN** state to the **SYN-RCVD** state

Target Name	Job Name	Metric	Description
		node_netstat_Tcp_CurrEstab	Number of TCP connections in the **ESTABLISHED** or **CLOSE-WAIT** state
		node_vmstat_pgmajfault	Number of major page faults in **vmstat**
		node_vmstat_pgpgout	Number of page out in **vmstat**
		node_vmstat_pgfault	Number of page faults in **vmstat**
		node_vmstat_pgpgin	Number of page in in **vmstat**
		node_processes_max_processes	Maximum number of processes
		node_processes_pids	Number of PIDs
		node_nf_conntrack_entries	Number of currently allocated flow entries for connection tracking
		node_nf_conntrack_entries_limit	Maximum size of a connection tracking table
		promhttp_metric_handler_requests_in_flight	Number of metrics being processed
		go_goroutines	Number of goroutines that exist
		node_filesystem_files	Number of files in the file system on the node
		node_filesystem_files_free	Number of available files in the file system on the node
podMonitor/monitoring/nvidia-gpu-device-plugin/0	monitoring/nvidia-gpu-device-plugin	cce_gpu_utilization	GPU compute usage
		cce_gpu_memory_utilization	GPU memory usage
		cce_gpu_encoder_utilization	GPU encoding usage
		cce_gpu_decoder_utilization	GPU decoding usage

Target Name	Job Name	Metric	Description
		cce_gpu_utilization_process	GPU compute usage of each process
		cce_gpu_memory_utilization_process	GPU memory usage of each process
		cce_gpu_encoder_utilization_process	GPU encoding usage of each process
		cce_gpu_decoder_utilization_process	GPU decoding usage of each process
		cce_gpu_memory_used	Used GPU memory
		cce_gpu_memory_total	Total GPU memory
		cce_gpu_memory_free	Free GPU memory
		cce_gpu_bar1_memory_used	Used GPU BAR1 memory
		cce_gpu_bar1_memory_total	Total GPU BAR1 memory
		cce_gpu_clock	GPU clock frequency
		cce_gpu_memory_clock	GPU memory frequency
		cce_gpu_graphics_clock	GPU frequency
		cce_gpu_video_clock	GPU video processor frequency
		cce_gpu_temperature	GPU temperature
		cce_gpu_power_usage	GPU power
		cce_gpu_total_energy_consumption	Total GPU energy consumption
		cce_gpu_pcie_link_bandwidth	GPU PCIe bandwidth
		cce_gpu_nvlink_bandwidth	GPU NVLink bandwidth
		cce_gpu_pcie_throughput_rx	GPU PCIe RX bandwidth
		cce_gpu_pcie_throughput_tx	GPU PCIe TX bandwidth

Target Name	Job Name	Metric	Description
		cce_gpu_nvlink_utilization_counter_rx	GPU NVLink RX bandwidth
		cce_gpu_nvlink_utilization_counter_tx	GPU NVLink TX bandwidth
		cce_gpu_retired_pages_sbe	Number of isolated GPU memory pages with single-bit errors
		cce_gpu_retired_pages_dbe	Number of isolated GPU memory pages with dual-bit errors
		xgpu_memory_total	Total xGPU memory
		xgpu_memory_used	Used xGPU memory
		xgpu_core_percentage_total	Total xGPU compute
		xgpu_core_percentage_used	Used xGPU compute
		gpu_schedule_policy	There are three GPU modes. **0**: GPU memory isolation, compute sharing mode. **1**: GPU memory and compute isolation mode. **2**: default mode, indicating that the GPU is not virtualized.
		xgpu_device_health	Health status of xGPU. **0**: xGPU is healthy. **1**: xGPU is unhealthy.
serviceMonitor/monitoring/prometheus-server/0	prometheus-server	prometheus_build_info	Prometheus build information
		prometheus_engine_query_duration_seconds	Time for query, in seconds
		prometheus_engine_query_duration_seconds_count	Number of queries
		prometheus_sd_discovered_targets	Number of metrics collected by different targets

Target Name	Job Name	Metric	Description
		prometheus_remote_storage_bytes_total	Total number of bytes of data (non-metadata) sent by the queue after compression
		prometheus_remote_storage_enqueue_retries_total	Number of retries upon enqueuing failed due to full shard queue
		prometheus_remote_storage_highest_timestamp_in_seconds	Latest timestamp in the remote storage
		prometheus_remote_storage_queue_highest_sent_timestamp_seconds	Highest timestamp successfully sent by remote storage
		prometheus_remote_storage_samples_dropped_total	Number of samples dropped before being sent to remote storage
		prometheus_remote_storage_samples_failed_total	Number of samples that failed to be sent to remote storage
		prometheus_remote_storage_samples_in_total	Number of samples sent to remote storage
		prometheus_remote_storage_samples_pending	Number of samples pending in shards to be sent to remote storage
		prometheus_remote_storage_samples_retried_total	Number of samples which failed to be sent to remote storage but were retried
		prometheus_remote_storage_samples_total	Total number of samples sent to remote storage
		prometheus_remote_storage_shard_capacity	Capacity of each shard of the queue used for parallel sending to the remote storage

Target Name	Job Name	Metric	Description
		prometheus_remote_storage_shards	Number of shards used for parallel sending to the remote storage
		prometheus_remote_storage_shards_desired	Number of shards that the queues shard calculation wants to run based on the rate of samples in vs. samples out
		prometheus_remote_storage_shards_max	Maximum number of shards that the queue is allowed to run
		prometheus_remote_storage_shards_min	Minimum number of shards that the queue is allowed to run
		prometheus_tsdb_wal_segment_current	WAL segment index that TSDB is currently writing to
		prometheus_tsdb_head_chunks	Number of chunks in the head block
		prometheus_tsdb_head_series	Number of time series stored in the head
		prometheus_tsdb_head_samples_appended_total	Number of appended samples
		prometheus_wal_watcher_current_segment	Current segment the WAL watcher is reading records from
		prometheus_target_interval_length_seconds	Metric collection interval
		prometheus_target_interval_length_seconds_count	Number of metric collection intervals
		prometheus_target_interval_length_seconds_sum	Sum of metric collection intervals
		prometheus_target_scrapes_exceeded_body_size_limit_total	Number of scrapes that hit the body size limit

Target Name	Job Name	Metric	Description
		prometheus_target_scrapes_exceeded_sample_limit_total	Number of scrapes that hit the sample limit
		prometheus_target_scrapes_sample_duplicate_timestamp_total	Number of scraped samples with duplicate timestamps
		prometheus_target_scrapes_sample_out_of_bounds_total	Number of samples rejected due to timestamp falling outside of the time bounds
		prometheus_target_scrapes_sample_out_of_order_total	Number of out-of-order samples
		prometheus_target_sync_length_seconds	Target synchronization interval
		prometheus_target_sync_length_seconds_count	Number of target synchronization intervals
		prometheus_target_sync_length_seconds_sum	Sum of target synchronization intervals
		promhttp_metric_handler_requests_in_flight	Current number of scrapes being served
		promhttp_metric_handler_requests_total	Total scrapes
		go_goroutines	Number of goroutines that exist
podMonitor/monitoring/virtual-kubelet-pods/0	monitoring/virtual-kubelet-pods	container_cpu_load_average_10s	Value of container CPU load average over the last 10 seconds
		container_cpu_system_seconds_total	Cumulative CPU time of a container system
		container_cpu_usage_seconds_total	Cumulative CPU time consumed by a container in core-seconds
		container_cpu_user_seconds_total	Cumulative CPU time of a user

Target Name	Job Name	Metric	Description
		container_cpu_cfs_periods_total	Number of elapsed enforcement period intervals
		container_cpu_cfs_throttled_periods_total	Number of throttled period intervals
		container_cpu_cfs_throttled_seconds_total	Total duration a container has been throttled
		container_fs_inodes_free	Number of available inodes in a file system
		container_fs_usage_bytes	File system usage
		container_fs_inodes_total	Number of inodes in a file system
		container_fs_io_current	Number of I/Os currently in progress in a disk or file system
		container_fs_io_time_seconds_total	Cumulative time spent on doing I/Os by the disk or file system
		container_fs_io_time_weighted_seconds_total	Cumulative weighted I/O time of a disk or file system
		container_fs_limit_bytes	Total disk or file system capacity that can be consumed by a container
		container_fs_reads_bytes_total	Cumulative amount of disk or file system data read by a container
		container_fs_read_seconds_total	Time a container spent on reading disk or file system data
		container_fs_reads_merged_total	Cumulative number of merged disk or file system reads made by a container

Target Name	Job Name	Metric	Description
		container_fs_reads_total	Cumulative number of disk or file system reads completed by a container
		container_fs_sector_reads_total	Cumulative number of disk or file system sector reads completed by a container
		container_fs_sector_writes_total	Cumulative number of disk or file system sector writes completed by a container
		container_fs_writes_bytes_total	Total amount of data written by a container to a disk or file system
		container_fs_write_seconds_total	Time a container spent on writing data to the disk or file system
		container_fs_writes_merged_total	Cumulative number of merged container writes to the disk or file system
		container_fs_writes_total	Cumulative number of disk or file system writes completed by a container
		container_blkio_device_usage_total	Blkio device bytes usage
		container_memory_failures_total	Cumulative number of container memory allocation failures
		container_memory_failcnt	Number of memory usage hits limits
		container_memory_cache	Memory used for the page cache of a container
		container_memory_mapped_file	Size of a container memory mapped file

Target Name	Job Name	Metric	Description
		container_memory_max_usage_bytes	Maximum memory usage recorded for a container
		container_memory_rss	Size of the resident memory set for a container
		container_memory_swap	Container swap usage
		container_memory_usage_bytes	Current memory usage of a container
		container_memory_working_set_bytes	Memory usage of the working set of a container
		container_network_receive_bytes_total	Total volume of data received by a container network
		container_network_receive_errors_total	Cumulative number of errors encountered during reception
		container_network_receive_packets_dropped_total	Cumulative number of packets dropped during reception
		container_network_receive_packets_total	Cumulative number of packets received
		container_network_transmit_bytes_total	Total volume of data transmitted on a container network
		container_network_transmit_errors_total	Cumulative number of errors encountered during transmission
		container_network_transmit_packets_dropped_total	Cumulative number of packets dropped during transmission
		container_network_transmit_packets_total	Cumulative number of packets transmitted
		container_processes	Number of processes running inside a container
		container_sockets	Number of open sockets for a container

Target Name	Job Name	Metric	Description
		container_file_descriptors	Number of open file descriptors for a container
		container_threads	Number of threads running inside a container
		container_threads_max	Maximum number of threads allowed inside a container
		container_ulimits_soft	Soft ulimit value of process 1 in a container Unlimited if the value is **-1**, except priority and nice.
		container_tasks_state	Number of tasks in the specified state, such as sleeping, running, stopped, uninterruptible, or ioawaiting
		container_spec_cpu_period	CPU period of a container
		container_spec_cpu_shares	CPU share of a container
		container_spec_cpu_quota	CPU quota of a container
		container_spec_memory_limit_bytes	Memory limit for a container
		container_spec_memory_reservation_limit_bytes	Memory reservation limit for a container
		container_spec_memory_swap_limit_bytes	Memory swap limit for a container
		container_start_time_seconds	Running time of a container
		container_last_seen	Last time a container was seen by the exporter
		container_accelerator_memory_used_bytes	GPU accelerator memory that is being used by a container

Target Name	Job Name	Metric	Description
		container_accelerator_memory_total_bytes	Total available memory of a GPU accelerator
		container_accelerator_duty_cycle	Percentage of time when a GPU accelerator is actually running
podMonitor/monitoring/everest-csi-controller/0	monitoring/everest-csi-controller	everest_action_result_total	Invoking of different functions
		everest_function_duration_seconds_bucket	Number of times that different functions are executed at different time
		everest_function_duration_seconds_count	Number of invoking times of different functions
		everest_function_duration_seconds_sum	Total invoking time of different functions
		everest_function_duration_quantile_seconds	Time quantile required for invoking different functions
		node_volume_read_completed_total	Number of completed reads
		node_volume_read_merged_total	Number of merged reads
		node_volume_read_bytes_total	Total number of bytes read by a sector
		node_volume_read_time_milliseconds_total	Total read duration
		node_volume_write_completed_total	Number of completed writes
		node_volume_write_merged_total	Number of merged writes
		node_volume_write_bytes_total	Total number of bytes written into a sector
		node_volume_write_time_milliseconds_total	Total write duration
		node_volume_io_now	Number of ongoing I/Os

Target Name	Job Name	Metric	Description
		node_volume_io_time_seconds_total	Total duration of I/O operations
		node_volume_capacity_bytes_available	Available capacity
		node_volume_capacity_bytes_total	Total capacity
		node_volume_capacity_bytes_used	Used capacity
		node_volume_inodes_available	Available inodes
		node_volume_inodes_total	Total number of inodes
		node_volume_inodes_used	Used inodes
		node_volume_read_transmissions_total	Number of read transmission times
		node_volume_read_timeouts_total	Number of read timeouts
		node_volume_read_sent_bytes_total	Number of bytes read
		node_volume_read_queue_time_milliseconds_total	Total read queue waiting time
		node_volume_read_rtt_time_milliseconds_total	Total read RTT
		node_volume_write_transmissions_total	Total number of write transmissions
		node_volume_write_timeouts_total	Total number of write timeouts
		node_volume_write_queue_time_milliseconds_total	Total write queue waiting time
		node_volume_write_rtt_time_milliseconds_total	Total write RTT
		node_volume_localvolume_stats_capacity_bytes	Total local volume capacity

Target Name	Job Name	Metric	Description
		node_volume_localvolume_stats_available_bytes	Available local volume capacity
		node_volume_localvolume_stats_used_bytes	Used local volume capacity
		node_volume_localvolume_stats_inodes	Number of inodes for a local volume
		node_volume_localvolume_stats_inodes_used	Used inodes for a local volume
podMonitor/monitoring/nginx-ingress-controller/0	monitoring/nginx-ingress-controller	nginx_ingress_controller_connect_duration_seconds_bucket	Duration for connecting to the upstream server
		nginx_ingress_controller_connect_duration_seconds_sum	Duration for connecting to the upstream server
		nginx_ingress_controller_connect_duration_seconds_count	Duration for connecting to the upstream server
		nginx_ingress_controller_request_duration_seconds_bucket	Time required for processing a request, in milliseconds
		nginx_ingress_controller_request_duration_seconds_sum	Time required for processing a request, in milliseconds
		nginx_ingress_controller_request_duration_seconds_count	Time required for processing a request, in milliseconds
		nginx_ingress_controller_request_size_bucket	Length of a request (including the request line, header, and body)
		nginx_ingress_controller_request_size_sum	Length of a request (including the request line, header, and body)
		nginx_ingress_controller_request_size_count	Length of a request (including the request line, header, and body)

Target Name	Job Name	Metric	Description
		nginx_ingress_controller_response_duration_seconds_bucket	Time required for receiving the response from the upstream server
		nginx_ingress_controller_response_duration_seconds_sum	Time required for receiving the response from the upstream server
		nginx_ingress_controller_response_duration_seconds_count	Time required for receiving the response from the upstream server
		nginx_ingress_controller_response_size_bucket	Length of a response (including the request line, header, and request body)
		nginx_ingress_controller_response_size_sum	Length of a response (including the request line, header, and request body)
		nginx_ingress_controller_response_size_count	Length of a response (including the request line, header, and request body)
		nginx_ingress_controller_header_duration_seconds_bucket	Time required for receiving the first header from the upstream server
		nginx_ingress_controller_header_duration_seconds_sum	Time required for receiving the first header from the upstream server
		nginx_ingress_controller_header_duration_seconds_count	Time required for receiving the first header from the upstream server
		nginx_ingress_controller_bytes_sent	Number of bytes sent to the client
		nginx_ingress_controller_ingress_upstream_latency_seconds	Upstream service latency
		nginx_ingress_controller_requests	Total number of client requests

Target Name	Job Name	Metric	Description
		nginx_ingress_controller_nginx_process_connections	Number of client connections in the active, read, write, or wait state
		nginx_ingress_controller_nginx_process_connections_total	Total number of client connections in the accepted or handled state
		nginx_ingress_controller_nginx_process_cpu_seconds_total	Total CPU time consumed by the Nginx process (unit: second)
		nginx_ingress_controller_nginx_process_num_procs	Number of processes
		nginx_ingress_controller_nginx_process_oldest_start_time_seconds	Start time in seconds since January 1, 1970
		nginx_ingress_controller_nginx_process_read_bytes_total	Total number of bytes read
		nginx_ingress_controller_nginx_process_requests_total	Total number of requests processed by Nginx since startup
		nginx_ingress_controller_nginx_process_resident_memory_bytes	Resident memory set usage of a process, that is, the actual physical memory usage
		nginx_ingress_controller_nginx_process_virtual_memory_bytes	Virtual memory usage of a process, that is, the total memory allocated to the process, including the actual physical memory and virtual swap space
		nginx_ingress_controller_nginx_process_write_bytes_total	Total amount of data written by the process to disks or other devices for long-term storage

Target Name	Job Name	Metric	Description
		nginx_ingress_controller_build_info	A metric with a constant '1' labeled with information about the build
		nginx_ingress_controller_check_success	Cumulative count of syntax check operations of the Nginx ingress controller
		nginx_ingress_controller_config_hash	Configured hash value
		nginx_ingress_controller_config_last_reload_successful	Whether the last configuration reload attempt was successful
		nginx_ingress_controller_config_last_reload_successful_timestamp_seconds	Timestamp of the last successful configuration reload
		nginx_ingress_controller_ssl_certificate_info	All information associated with a certificate
		nginx_ingress_controller_success	Cumulative number of reload operations of the Nginx ingress controller
		nginx_ingress_controller_orphan_ingress	Status of an orphaned ingress (**1** indicates an orphaned ingress). 0: Not isolated. **namespace**: namespace of the ingress **ingress**: name of the ingress **type**: status of the ingress. The value can be **no-service** or **no-endpoint**.
		nginx_ingress_controller_admission_config_size	Size of the admission controller configuration

Target Name	Job Name	Metric	Description
		nginx_ingress_controller_admission_render_duration	Rendering duration of the admission controller
		nginx_ingress_controller_admission_render_ingresses	Length of ingresses rendered by the admission controller
		nginx_ingress_controller_admission_roundtrip_duration	Time spent by the admission controller to process new events
		nginx_ingress_controller_admission_tested_duration	Time spent on admission controller tests
		nginx_ingress_controller_admission_tested_ingresses	Length of ingresses processed by the admission controller
podMonitor/monitoring/cceaddon-npd/0	monitoring/cceaddon-npd	problem_counter	Number of times that the check item is found abnormal
		problem_gauge	Whether the check item has triggered an exception • 0: not triggered • 1: triggered

10.4 Basic Metrics: ModelArts Metrics

This section describes the ModelArts metrics reported to AOM through the Agent.

Table 10-4 Metrics reported by ModelArts to AOM through the Agent

Category	Metric	Metric Name	Description	Value Range	Unit
CPU	ma_container_cpu_util	CPU Usage	CPU usage of a measured object	0–100	%
	ma_container_cpu_used_core	Used CPU Cores	Number of CPU cores used by a measured object	≥ 0	Cores

Category	Metric	Metric Name	Description	Value Range	Unit
	ma_container_cpu_limit_core	Total CPU Cores	Total number of CPU cores that have been applied for a measured object	≥ 1	Cores
Memory	ma_container_memory_capacity_megabytes	Memory	Total physical memory that has been applied for a measured object	≥ 0	MB
	ma_container_memory_util	Physical Memory Usage	Percentage of the used physical memory to the total physical memory applied for a measured object	0–100	%
	ma_container_memory_used_megabytes	Used Physical Memory	Physical memory that has been used by a measured object (**container_memory_working_set_bytes** in the current working set). (Memory usage in a working set = Active anonymous and cache, and file-baked page ≤ **container_memory_usage_bytes**)	≥ 0	MB
Storage I/O	ma_container_disk_read_kilobytes	Disk Read Rate	Volume of data read from a disk per second	≥ 0	KB/s
	ma_container_disk_write_kilobytes	Disk Write Rate	Volume of data written into a disk per second	≥ 0	KB/s
GPU memory	ma_container_gpu_mem_total_megabytes	GPU Memory Capacity	Total GPU memory of a training job	> 0	MB
	ma_container_gpu_mem_util	GPU Memory Usage	Percentage of the used GPU memory to the total GPU memory	0–100	%

Category	Metric	Metric Name	Description	Value Range	Unit
	ma_container_gpu_mem_used_megabytes	Used GPU Memory	GPU memory used by a measured object	≥ 0	MB
GPU	ma_container_gpu_util	GPU Usage	GPU usage of a measured object	0–100	%
	ma_container_gpu_mem_copy_util	GPU Memory Bandwidth Usage	GPU memory bandwidth usage of a measured object. For example, the maximum memory bandwidth of NVIDIA GPU V100 is 900 GB/s. If the current memory bandwidth is 450 GB/s, the memory bandwidth usage is 50%.	0–100	%
	ma_container_gpu_enc_util	GPU Encoder Usage	GPU encoder usage of a measured object	0–100	%
	ma_container_gpu_dec_util	GPU Decoder Usage	GPU decoder usage of a measured object	0–100	%
	DCGM_FI_DEV_GPU_TEMP	GPU Temperature	GPU temperature	> 0	°C
	DCGM_FI_DEV_POWER_USAGE	GPU Power	GPU power	> 0	W
	DCGM_FI_DEV_MEMORY_TEMP	Memory Temperature	Memory temperature	> 0	°C

Category	Metric	Metric Name	Description	Value Range	Unit
	DCGM_FI_PROF_GR_ENGINE_ACTIVE	Graphics Engine Activity	Percentage of the time when the graphic or compute engine is in the active state within a period. This is an average value of all graphic or compute engines. An active graphic or compute engine indicates that the graphic or compute context is associated with a thread and the graphic or compute context is busy.	0–1.0	Percentage (fraction)
	DCGM_FI_PROF_SM_OCCUPANCY	SM Occupancy	Ratio of the number of thread bundles that reside on the SM to the maximum number of thread bundles that can reside on the SM within a period. This is an average value of all SMs within a period. A high value does not mean a high GPU usage. Only when the GPU memory bandwidth is limited, a high value of workloads (**DCGM_FI_PROF_DRAM_ACTIVE**) indicates more efficient GPU usage.	0–1.0	Percentage (fraction)

Category	Metric	Metric Name	Description	Value Range	Unit
	DCGM_FI_PROF_PIPE_TENSOR_ACTIVE	Tensor Activity	Fraction of the period during which the tensor (HMMA/IMMA) pipe is active. This is an average value within a period, not an instantaneous value. A higher value indicates a higher utilization of tensor cores. Value 1 (100%) indicates that a tensor instruction is sent every instruction cycle in the entire period (one instruction is completed in two cycles). If the value is 0.2 (20%), the possible causes are as follows: During the entire period, 20% of the SM tensor cores run at 100% utilization. During the entire period, all SM tensor cores run at 20% utilization. During 1/5 of the entire period, all SM tensor cores run at 100% utilization. Other combinations	0–1.0	Percentage (fraction)

Category	Metric	Metric Name	Description	Value Range	Unit
	DCGM_FI_PROF_DRAM_ACTIVE	Memory BW Utilization	Percentage of the time for sending data to or receiving data from the device memory within a period. This is an average value within a period, not an instantaneous value. A higher value indicates a higher utilization of device memory. Value 1 (100%) indicates that a DRAM instruction is executed once per cycle throughout a period (the maximum value can be reached at a peak of about 0.8). If the value is 0.2 (20%), indicating that data is read from or written into the device memory during 20% of the cycle within a period.	0–1.0	Percentage (fraction)

Category	Metric	Metric Name	Description	Value Range	Unit
	DCGM_FI_PROF_PIPE_FP16_ACTIVE	FP16 Engine Activity	Fraction of the period during which the FP16 (half-precision) pipe is active. This is an average value within a period, not an instantaneous value. A larger value indicates a higher usage of FP16 cores. Value 1 (100%) indicates that the FP16 instruction is executed every two cycles (for example, Volta cards) in a period. If the value is 0.2 (20%), the possible causes are as follows: During the entire period, 20% of the SM FP16 cores run at 100% utilization. During the entire period, all SM FP16 cores run at 20% utilization. During 1/5 of the entire period, all SM FP16 cores run at 100% utilization. Other combinations	0–1.0	Percentage (fraction)

Category	Metric	Metric Name	Description	Value Range	Unit
	DCGM_FI_PROF_PIPE_FP32_ACTIVE	FP32 Engine Activity	Fraction of the period during which the fused multiply-add (FMA) pipe is active. Multiply-add applies to FP32 (single precision) and integers. This is an average value within a period, not an instantaneous value. A larger value indicates a higher usage of FP32 cores. Value 1 (100%) indicates that the FP32 instruction is executed every two cycles (for example, Volta cards) in a period. If the value is 0.2 (20%), the possible causes are as follows: During the entire period, 20% of the SM FP32 cores run at 100% utilization. During the entire period, all SM FP32 cores run at 20% utilization. During 1/5 of the entire period, all SM FP32 cores run at 100% utilization. Other combinations	0–1.0	Percentage (fraction)

Category	Metric	Metric Name	Description	Value Range	Unit
	DCGM_FI_PROF_PIPE_FP64_ACTIVE	FP64 Engine Activity	Fraction of the period during which the FP64 (double precision) pipe is active. This is an average value within a period, not an instantaneous value. A larger value indicates a higher usage of FP64 cores. Value 1 (100%) indicates that the FP64 instruction is executed every four cycles (for example, Volta cards) in a period. If the value is 0.2 (20%), the possible causes are as follows: During the entire period, 20% of the SM FP64 cores run at 100% utilization. During the entire period, all SM FP64 cores run at 20% utilization. During 1/5 of the entire period, all SM FP64 cores run at 100% utilization. Other combinations	0–1.0	Percentage (fraction)

Category	Metric	Metric Name	Description	Value Range	Unit
	DCGM_FI_PROF_SM_ACTIVE	SM Activity	Fraction of the time during which at least one thread bundle is active on an SM within a period. This is an average value of all SMs and is insensitive to the number of threads in each block. A thread bundle is active after being scheduled and allocated with resources. The thread bundle may be in the computing state or a non-computing state (for example, waiting for a memory request). If the value is less than 0.5, GPUs are not efficiently used. The value should be greater than 0.8. For example, a GPU has N SMs: A kernel function uses N thread blocks to run on all SMs in a period. In this case, the value is 1 (100%). A kernel function runs N/5 thread blocks in a period. In this case, the value is 0.2. A kernel function uses N thread blocks and runs only 1/5 of cycles in a period. In this case, the value is 0.2.	0–1.0	Percentage (fraction)

Category	Metric	Metric Name	Description	Value Range	Unit
	DCGM_FI_PROF_PCIE_TX_BYTES DCGM_FI_PROF_PCIE_RX_BYTES	PCIe Bandwidth	Rate of data transmitted or received over the PCIe bus, including the protocol header and data payload. This is an average value within a period, not an instantaneous value. The rate is averaged over the period. For example, if 1 GB of data is transmitted within 1 second, the transmission rate is 1 GB/s regardless of whether the data is transmitted at a constant rate or burst. Theoretically, the maximum PCIe Gen3 bandwidth is 985 MB/s per channel.	≥ 0	Bytes/s

Category	Metric	Metric Name	Description	Value Range	Unit
	DCGM_FI_PROF_NVLINK_RX_BYTES DCGM_FI_PROF_NVLINK_TX_BYTES	NVLink Bandwidth	Rate at which data is transmitted or received through NVLink, excluding the protocol header. This is an average value within a period, not an instantaneous value. The rate is averaged over the period. For example, if 1 GB of data is transmitted within 1 second, the transmission rate is 1 GB/s regardless of whether the data is transmitted at a constant rate or burst. Theoretically, the maximum NVLink Gen2 bandwidth is 25 GB/s per link in each direction.	≥ 0	Bytes/s
Network I/O	ma_container_network_receive_bytes	Downlink Rate (BPS)	Inbound traffic rate of a measured object	≥ 0	Bytes/s
	ma_container_network_receive_packets	Downlink Rate (PPS)	Number of data packets received by a NIC per second	≥ 0	Packets/s
	ma_container_network_receive_error_packets	Downlink Error Rate	Number of error packets received by a NIC per second	≥ 0	Count/s
	ma_container_network_transmit_bytes	Uplink Rate (BPS)	Outbound traffic rate of a measured object	≥ 0	Bytes/s

Category	Metric	Metric Name	Description	Value Range	Unit
	ma_container_network_transmit_error_packets	Uplink Error Rate	Number of error packets sent by a NIC per second	≥ 0	Count/s
	ma_container_network_transmit_packets	Uplink Rate (PPS)	Number of data packets sent by a NIC per second	≥ 0	Packets/s
NPU	ma_container_npu_util	NPU Usage	NPU usage of a measured object	0–100	%
	ma_container_npu_memory_util	NPU Memory Usage	Percentage of the used NPU memory to the total NPU memory	0–100	%
	ma_container_npu_memory_used_megabytes	Used NPU Memory	NPU memory used by a measured object	≥ 0	MB
	ma_container_npu_memory_total_megabytes	Total NPU Memory	Total NPU memory of a measured object	≥ 0	MB

10.5 Basic Metrics: IEF Metrics

This section describes the types, names, and meanings of IEF metrics reported to AOM. **After IEF metrics are reported to AOM, AOM will convert them based on mapping rules and displays the results on the Metric Browsing page.**

Table 10-5 IEF metrics

Category	Sub-Category	Metrics Displayed on AOM	Metrics Reported by IEF	Metric Name	Description	Value Range	Unit
Host metrics	CPU	aom_node_cpu_limit_core	cpuCoreLimit	Total CPU Cores	Total number of CPU cores that have been applied for a measured object	≥ 1	Cores
		aom_node_cpu_used_core	cpuCoreUsed	Used CPU Cores	Number of CPU cores used by a measured object	≥ 0	Cores
		aom_node_cpu_usage	cpuUsage	CPU Usage	CPU usage of a measured object	0–100	%
	Memory	aom_node_memory_total_megabytes	totalMem	Total Physical Memory	Total physical memory that has been applied for a measured object	≥ 0	MB
		aom_node_memory_free_megabytes	freeMem	Available Physical Memory	Available physical memory of a measured object	≥ 0	MB
		aom_node_memory_usage	memUsedRate	Physical Memory Usage	Percentage of the used physical memory to the total physical memory applied for a measured object	0–100	%
		aom_node_virtual_memory_usage	virMemUsedRate	Virtual Memory Usage	Percentage of the used virtual memory to the total virtual memory	≥ 0	%

Category	Sub-Category	Metrics Displayed on AOM	Metrics Reported by IEF	Metric Name	Description	Value Range	Unit
	Network	aom_node_network_receive_bytes	recvBytesRate	Downlink Rate (BPS)	Inbound traffic rate of a measured object	≥ 0	Bytes/s
		aom_node_network_transmit_bytes	sendBytesRate	Uplink Rate (BPS)	Outbound traffic rate of a measured object	≥ 0	Bytes/s
	Disk	aom_node_disk_capacity_megabytes	diskCapacity	Total Disk Space	Total disk space	≥ 0	MB
		aom_node_disk_available_capacity_megabytes	diskAvailableCapacity	Available Disk Space	Disk space that has not been used	≥ 0	MB
		aom_node_disk_usage	diskUsedRate	Disk Usage	Percentage of the used disk space to the total disk space	0–100	%
		aom_node_disk_read_kilobytes	diskReadRate	Disk Read Rate	Volume of data read from a disk per second	≥ 0	KB/s
		aom_node_disk_write_kilobytes	diskWriteRate	Disk Write Rate	Volume of data written into a disk per second	≥ 0	KB/s
	GPU	aom_node_gpu_memory_free_megabytes	gpuMemCapacity	GPU Memory	GPU memory of a measured object	≥ 0	MB

Category	Sub-Category	Metrics Displayed on AOM	Metrics Reported by IEF	Metric Name	Description	Value Range	Unit
		aom_node_gpu_memory_usage	gpuMemUsage	GPU Memory Usage	Percentage of the used GPU memory to the total GPU memory	0–100	%
		aom_node_gpu_memory_used_megabytes	gpuMemUsed	Used GPU Memory	GPU memory used by a measured object	≥ 0	MB
		aom_node_gpu_usage	gpuUtil	GPU Usage	GPU usage of a measured object	0–100	%
	Host	aom_node_process_number	processNum	Processes	Number of running processes on a measured object	≥ 0	N/A
	Atlas 500 AI Edge Station	aom_node_npu_temperature_centigrade	node_temperature	Node Temperature	Temperature of the Atlas 500 AI Edge Station node, which is reported by calling the edgecore API	≥ 0	°C
		node_power	node_power	Node Power	Power of the Atlas 500 AI Edge Station node, which is reported by calling the edgecore API	≥ 0	W
		node_voltage	node_voltage	Node Voltage	Voltage of the Atlas 500 AI Edge Station node, which is reported by calling the edgecore API	≥ 0	V
		npu_temperature	npu_temperature	Chip Temperature	NPU temperature of the Atlas 500 AI Edge Station node, which is reported by calling the edgecore API	≥ 0	°C

Category	Sub-Category	Metrics Displayed on AOM	Metrics Reported by IEF	Metric Name	Description	Value Range	Unit
		npu_health	npu_health	Chip Health Status	NPU health status of the Atlas 500 AI Edge Station node, which is reported by calling the edgecore API	≥ 0	N/A
		ai_cpu_rate	ai_cpu_rate	AI CPU Usage	AI CPU usage of the Ascend AI accelerator card, which is reported by calling the edgecore API	0–100	%
		ai_core_rate	ai_core_rate	AI Core Usage	AI Core usage of the Ascend AI accelerator card, which is reported by calling the edgecore API	0–100	%
		ctrl_cpu_rate	ctrl_cpu_rate	Control CPU Usage	Control CPU usage of the Ascend AI accelerator card, which is reported by calling the edgecore API	0–100	%
		ddr_cap_rate	ddr_cap_rate	DDR Memory Usage	DDR memory usage of the Atlas 500 AI Edge Station node, which is reported by calling the edgecore API	0–100	%
		ddr_bw_rate	ddr_bw_rate	DDR Bandwidth Usage	DDR bandwidth usage of the Atlas 500 AI Edge Station node, which is reported by calling the edgecore API	0–100	%
Container metrics	CPU	aom_container_cpu_limit_core	cpuCoreLimit	Total CPU Cores	Total number of CPU cores that have been applied for a measured object	≥ 1	Cores

Category	Sub-Category	Metrics Displayed on AOM	Metrics Reported by IEF	Metric Name	Description	Value Range	Unit
		aom_container_cpu_used_core	cpuCoreUsed	Used CPU Cores	Number of CPU cores used by a measured object	≥ 0	Cores
		aom_container_cpu_usage	cpuUsage	CPU Usage	CPU usage of a measured object	0–100	%
	Memory	aom_container_memory_request_megabytes	memCapacity	Total Physical Memory	Total physical memory that has been applied for a measured object	≥ 0	MB
		aom_container_memory_used_megabytes	memUsed	Used Physical Memory	Used physical memory of a measured object	≥ 0	MB
		memUsedRate	memUsedRate	Physical Memory Usage	Percentage of the used physical memory to the total physical memory applied for a measured object	0–100	%
	Disk	aom_container_disk_read_kilobytes	diskReadRate	Disk Read Rate	Volume of data read from a disk per second	≥ 0	KB/s
		aom_container_disk_write_kilobytes	diskWriteRate	Disk Write Rate	Volume of data written into a disk per second	≥ 0	KB/s
	Network	aom_container_network_receive_bytes	recvBytesRate	Downlink Rate (BPS)	Inbound traffic rate of a measured object	≥ 0	Bytes/s

Category	Sub-Category	Metrics Displayed on AOM	Metrics Reported by IEF	Metric Name	Description	Value Range	Unit
		aom_container_network_transmit_bytes	sendBytesRate	Uplink Rate (BPS)	Outbound traffic rate of a measured object	≥ 0	Bytes/s
	GPU	aom_container_gpu_memory_free_megabytes	gpuMemCapacity	GPU Memory	GPU memory of a measured object	≥ 0	MB
		aom_container_gpu_memory_usage	gpuMemUsage	GPU Memory Usage	Percentage of the used GPU memory to the total GPU memory	0–100	%
		aom_container_gpu_memory_used_megabytes	gpuMemUsed	Used GPU Memory	GPU memory used by a measured object	≥ 0	MB
		aom_container_gpu_usage	gpuUtil	GPU Usage	GPU usage of a measured object	0–100	%
	Container status	aom_container_status	status	Container Status	Container status	≥ 0	N/A
Process metrics	CPU	aom_process_cpu_usage	cpuUsage	CPU Usage	CPU usage of a measured object	0–100	%
	Memory	aom_process_memory_used_megabytes	memUsed	Used Physical Memory	Used physical memory of a measured object	≥ 0	MB

Category	Sub-Category	Metrics Displayed on AOM	Metrics Reported by IEF	Metric Name	Description	Value Range	Unit
	Process status	aom_process_status	status	Process Status	Process Status	≥ 0	N/A
	GPU	gpuMemCapacity	gpuMemCapacity	GPU Memory	GPU memory of a measured object	≥ 0	MB
		gpuMemUsage	gpuMemUsage	GPU Memory Usage	Percentage of the used GPU memory to the total GPU memory	0–100	%
		gpuMemUsed	gpuMemUsed	Used GPU Memory	GPU memory used by a measured object	≥ 0	MB
		gpuUtil	gpuUtil	GPU Usage	GPU usage of a measured object	0–100	%

10.6 Basic Metrics: CSE Metrics

This section describes the types, names, and meanings of Cloud Service Engine (CSE) metrics reported to AOM.

Table 10-6 CSE metrics

Category	Sub-Category	Metric	Metric Name	Description	Value Range	Unit
ServiceComb	registry	servicecomb_service_center_db_service_total	Microservice Versions	Number of microservice versions	≥ 0	Count

Category	Sub-Category	Metric	Metric Name	Description	Value Range	Unit
		servicecomb_service_center_db_instance_total	Microservice Instances	Number of microservice instances	≥ 0	Count
		servicecomb_service_center_http_request_total	HTTP Requests	Number of HTTP requests, covering multiple URLs, methods, and codes	≥ 0	Count
		servicecomb_service_center_http_request_durations_microseconds	Total HTTP Request Time	Total HTTP request time, covering multiple URLs, methods, and codes	≥ 0	μs
	config	servicecomb_kie_request_count	HTTP Requests	Number of HTTP requests, covering multiple URLs, methods, and codes	≥ 0	Count
		servicecomb_kie_request_process_duration	Total HTTP Request Time	Total HTTP request time, covering multiple URLs, methods, and codes	≥ 0	ms
		servicecomb_kie_config_count	Configs	Number of ServiceComb configs	≥ 0	Count
Nacos	config	configCount	Nacos Configs	Number of configs in each Nacos cluster node	≥ 0	Count

Category	Sub-Category	Metric	Metric Name	Description	Value Range	Unit
		longPolling	HTTP Persistent Connections of Nacos Config (Listeners)	Number of HTTP persistent connections of Nacos config	≥ 0	Count
	http	http_server_requests_seconds_count	HTTP Requests	Number of HTTP requests, covering multiple URLs, methods, and codes	≥ 0	Count
		http_server_requests_seconds_max	Maximum HTTP Request Time	Maximum HTTP request time, covering URLs, methods, and codes. This parameter is reported when Nacos-Client 1.x is used.	≥ 0	s
		http_server_requests_seconds_sum	Total HTTP Request Time	Total HTTP request time, covering multiple URLs, methods, and codes	≥ 0	s
	naming	avgPushCost	Avg. Nacos Naming Push Time	Average Nacos naming push time	≥ 0	ms
		maxPushCost	Max. Nacos Naming Push Time	Maximum Nacos naming push time	≥ 0	ms

Category	Sub-Category	Metric	Metric Name	Description	Value Range	Unit
		ipCount	Nacos Naming IP Addresses	Number of microservice instances that are registered	≥ 0	Count
		subscriberCount	Nacos Naming Subscribers	Number of Nacos naming subscribers	≥ 0	Count
		serviceCount	Nacos Naming Domain Names (2.x)	Number of services in each Nacos cluster node	≥ 0	Count
	instance	jvm_memory_max_bytes	Max. JVM Memory	Maximum memory size, including the memory of various areas	≥ 0	Bytes
		jvm_memory_used_bytes	Used JVM Memory	Used memory, including the memory of various areas	≥ 0	Bytes
		up	Instance Running Status	Whether the Nacos instance is running • **1**: running • **0**: not running	0 or 1	N/A
Application gateway	envoy	cpuUsage	CPU Usage	CPU usage of a measured object	0–100	%
		envoy_http_downstream_cx_active	Active Connections	Number of active connections	≥ 0	Count
		downstream_cx_delayed_close_timeout	Connections Delayed to Close	Number of connections that are delayed to close	≥ 0	Count

Category	Sub-Category	Metric	Metric Name	Description	Value Range	Unit
		envoy_http_downstream_cx_destroy	Destroyed Connections	Number of connections that are destroyed	≥ 0	Count
		envoy_http_downstream_cx_destroy_active_rq	Destroyed Active Connections	Number of active connections that are destroyed	≥ 0	Count
		envoy_http_downstream_cx_destroy_local	Destroyed Local Connections	Number of local connections that are destroyed	≥ 0	Count
		envoy_http_downstream_cx_destroy_local_active_rq	Destroyed Local Active Connections	Number of local active connections that are destroyed	≥ 0	Count
		envoy_http_downstream_cx_destroy_remote	Destroyed Connections Due to Remote Shutdown	Number of connections that are destroyed due to remote shutdown	≥ 0	Count
		envoy_http_downstream_cx_destroy_remote_active_rq	Destroyed Active Connections Due to Remote Shutdown	Number of active connections that are destroyed due to remote shutdown	≥ 0	Count
		envoy_http_downstream_cx_drain_close	Closed Connections Due to Ejections	Number of connections that are closed due to ejections	≥ 0	Count

Category	Sub-Category	Metric	Metric Name	Description	Value Range	Unit
		envoy_http_downstream_cx_http1_active	HTTP1 Connections	Number of HTTP1 connections	≥ 0	Count
		envoy_http_downstream_cx_max_duration_reached	Timeout Connections	Number of connections that timed out	≥ 0	Count
		envoy_http_downstream_cx_tx_bytes_total	Total Sent Bytes	Total number of bytes that are sent	≥ 0	Bytes
		envoy_http_downstream_rq	Total Requests	Total number of requests	≥ 0	Count
		envoy_http_downstream_rq_http1_total	Total HTTP1 Requests	Total number of HTTP1 requests	≥ 0	Count
		envoy_http_downstream_rq_http2_total	Total HTTP2 Requests	Total number of HTTP2 requests	≥ 0	Count
		envoy_http_downstream_rq_idle_timeout	Closed Requests Due to Excessive Idle Time	Number of requests that are closed due to excessive idle time	≥ 0	Count
		envoy_http_downstream_rq_too_large	Requests with Too Large Bodies	Number of requests with too large bodies (status code **413** returned)	≥ 0	Count
		downstream_rq_ws_on_non_ws_route	WebSocket Requests Without Routes	Number of requests that are rejected due to a lack of routes	≥ 0	Count

Category	Sub-Category	Metric	Metric Name	Description	Value Range	Unit
		envoy_http_local_rate_limiter_http_local_rate_limit_enforced	Limited Requests	Number of requests that are limited	≥ 0	Count
		envoy_cluster_circuit_breakers_default_cx_open	Connection Circuit Breaker Trigger Status	0: The concurrency limit has not been reached. 1: The concurrency limit has been reached. No more requests can be accepted.	0 or 1	N/A
		envoy_cluster_circuit_breakers_high_cx_open				
		envoy_cluster_circuit_breakers_default_cx_pool_open	Pool's Circuit Breaker Trigger Status	0: The concurrency limit has not been reached. 1: The concurrency limit has been reached. No more requests can be accepted.	0 or 1	N/A
		envoy_cluster_circuit_breakers_high_cx_pool_open				
		envoy_cluster_circuit_breakers_default_remaining_cx	Remaining Connections	Number of remaining connections that can be accepted by the connection circuit breaker	≥ 0	Count
		envoy_cluster_circuit_breakers_high_remaining_cx				
		envoy_cluster_circuit_breakers_default_remaining_cx_pools	Pool's Remaining Connections	Number of remaining connections that can be accepted by the pool circuit breaker	≥ 0	Count

Category	Sub-Category	Metric	Metric Name	Description	Value Range	Unit
		envoy_cluster_circuit_breakers_high_remaining_cx_pools				
		envoy_cluster_circuit_breakers_default_remaining_pending	Pending Requests	Number of requests to be processed before the circuit breaker reaches the concurrency limit	≥ 0	Count
		envoy_cluster_circuit_breakers_high_remaining_pending				
		envoy_cluster_circuit_breakers_default_remaining_retries	Remaining Retries	Number of remaining retries before the circuit breaker reaches the concurrency limit	≥ 0	Count
		envoy_cluster_circuit_breakers_high_remaining_retries				
		envoy_cluster_circuit_breakers_default_remaining_rq	Remaining Requests	Number of remaining requests before the circuit breaker reaches the concurrency limit	≥ 0	Count
		envoy_cluster_circuit_breakers_high_remaining_rq				

Category	Sub-Category	Metric	Metric Name	Description	Value Range	Unit
		envoy_cluster_circuit_breakers_default_rq_open envoy_cluster_circuit_breakers_high_rq_open	Request Circuit Breaker Trigger Status	0: The concurrency limit has not been reached. 1: The concurrency limit has been reached. No more requests can be accepted.	0 or 1	N/A
		envoy_cluster_circuit_breakers_default_rq_retry_open envoy_cluster_circuit_breakers_high_rq_retry_open	Retry Circuit Breaker Trigger Status	0: The concurrency limit has not been reached. 1: The concurrency limit has been reached. No more requests can be accepted.	0 or 1	N/A
		envoy_cluster_ejections_overflow	Ejections Due to Overflow	Number of ejections occurred due to overflow	≥ 0	Count
		envoy_cluster_ejections_consecutive_5xx	Ejections Caused by Consecutive 5xx Errors	Number of ejections that are caused by consecutive 5xx errors	≥ 0	Count
		envoy_cluster_ejections_detected_consecutive_5xx	Detected Ejections Caused by Consecutive 5xx Errors	Number of detected ejections (even if not forcibly enforced) that are caused by consecutive 5xx errors	≥ 0	Count

Category	Sub-Category	Metric	Metric Name	Description	Value Range	Unit
		envoy_cluster_ejections_detected_consecutive_gateway_failure	Detected Ejections Caused by Consecutive Gateway Faults	Number of detected ejections (even if not forcibly enforced) that are caused by consecutive gateway faults	≥ 0	Count
		envoy_cluster_ejections_detected_consecutive_local_origin_failure	Detected Ejections Caused by Consecutive Local Faults	Number of detected ejections (even if not forcibly enforced) that are caused by consecutive local faults	≥ 0	Count
		envoy_cluster_ejections_enforced_consecutive_local_origin_failure	Forced Ejections Caused by Consecutive Local Faults	Number of forced ejections that are caused by consecutive local faults	≥ 0	Count
		envoy_cluster_ejections_detected_failure_percentage	Ejections Caused by High Request Failure Rate	Number of ejections occurred because the request failure rate exceeds the threshold	≥ 0	Count
		envoy_cluster_ejections_detected_local_origin_failure_percentage	Detected ejections Caused by High Local Request Failure Rate	Number of ejections occurred because the local request failure rate exceeds the threshold	≥ 0	Count

Category	Sub-Category	Metric	Metric Name	Description	Value Range	Unit
		envoy_cluster_ejections_detected_local_origin_success_rate	Detected Ejections Caused by Low Local Request Success Rate	Number of ejections occurred (even if not forcibly enforced) because the local request success rate does not reach the threshold	≥ 0	Count
		envoy_cluster_ejections_detected_success_rate	Detected Ejections Caused by Low Local Request Success Rate	Number of ejections occurred because the request success rate does not reach the threshold	≥ 0	Count
		envoy_cluster_ejections_enforced_consecutive_5xx	Enforced Ejections Caused by Consecutive 5xx Errors	Number of forced ejections that are caused by consecutive 5xx errors	≥ 0	Count
		envoy_cluster_ejections_enforced_consecutive_gateway_failure	Forced Ejections Caused by Consecutive Gateway Faults	Number of forced ejections that are caused by consecutive gateway faults	≥ 0	Count
		envoy_cluster_ejections_enforced_failure_percentage	Forced Ejections Caused by High Request Failure Rate	Number of forced ejections occurred because the request failure rate exceeds the threshold	≥ 0	Count

Category	Sub-Category	Metric	Metric Name	Description	Value Range	Unit
		envoy_cluster_ejections_enforced_local_origin_failure_percentage	Forced Ejections Caused by High Local Request Failure Rate	Number of forced ejections occurred because the local request failure rate exceeds the threshold	≥ 0	Count
		envoy_cluster_ejections_enforced_local_origin_success_rate	Forced Ejections Caused by Low Local Request Success Rate	Number of forced ejections occurred because the local request success rate does not reach the threshold	≥ 0	Count
		envoy_cluster_ejections_enforced_success_rate	Forced Ejections Caused by Low Request Success Rate	Number of forced ejections occurred because the request success rate does not reach the threshold	≥ 0	Count
		envoy_cluster_ejections_enforced_total	Forced Ejections	Number of forced ejections that are caused by any exception	≥ 0	Count
		envoy_http_downstream_cx_rx_bytes_total	Total Bytes Received	Total number of bytes that are received	≥ 0	Bytes

10.7 Basic Metrics: IoTDA Metrics

This section describes the types, names, and meanings of IoTDA metrics reported to AOM.

Table 10-7 IoTDA metrics

Category	Metric	Metric Name	Description	Value Range	Unit
Total devices	iotda_device_status_onlineCount	Online Devices	Number of online devices of a tenant in the current instance or resource space	≥ 0	Count
	iotda_device_status_totalCount	Total Devices	Total number of devices of a tenant in the current instance or resource space	≥ 0	Count
	iotda_device_status_offlineCount	Offline Devices	Number of offline devices of a tenant in the current instance or resource space	≥ 0	Count
	iotda_device_status_abnormalCount	Abnormal Devices	Number of abnormal devices of a tenant in the current instance or resource space	≥ 0	Count
	iotda_device_status_inactiveCount	Inactive Devices	Number of inactive devices of a tenant in the current instance or resource space	≥ 0	Count
	iotda_device_status_activeCount	Active Devices	Number of active devices of a tenant in the current instance or resource space	≥ 0	Count
	iotda_device_status_dailyOnlineCount	Online Devices (Accumulated)	Accumulated number of online devices of a tenant in the current instance or resource space	≥ 0	Count
Reported NB-IoT data records	iotda_south_dataReport_totalCount	Total NB-IoT Data Reporting Times	Total number of times NB-IoT data is reported in the current instance or resource space	≥ 0	Count
	iotda_south_dataReport_successCount	Successful NB-IoT Data Reporting Times	Number of times NB-IoT data is successfully reported in the current instance or resource space	≥ 0	Count

Category	Metric	Metric Name	Description	Value Range	Unit
	iotda_south_dataReport_failedCount	Failed NB-IoT Data Reporting Times	Number of times NB-IoT data fails to be reported in the current instance or resource space	≥ 0	Count
MQTT event reporting times	iotda_south_eventUp_totalCount	Total MQTT Event Reporting Times	Total number of times MQTT events are reported in the current instance or resource space	≥ 0	Count
	iotda_south_eventUp_successCount	Successful MQTT Event Reporting Times	Number of times MQTT events are successfully reported in the current instance or resource space	≥ 0	Count
	iotda_south_eventUp_failedCount	Failed MQTT Event Reporting Times	Number of times MQTT events fail to be reported in the current instance or resource space	≥ 0	Count
MQTT property reporting times	iotda_south_propertiesReport_totalCount	Total MQTT Property Reporting Times	Total number of times MQTT properties are reported in the current instance or resource space	≥ 0	Count
	iotda_south_propertiesReport_successCount	Successful MQTT Property Reporting Times	Number of times MQTT properties are successfully reported in the current instance or resource space	≥ 0	Count
	iotda_south_propertiesReport_failedCount	Failed MQTT Property Reporting Times	Number of times MQTT properties fail to be reported in the current instance or resource space	≥ 0	Count

Category	Metric	Metric Name	Description	Value Range	Unit
MQTT message reporting times	iotda_south_messageUp_totalCount	Total MQTT Message Reporting Times	Total number of times MQTT messages are reported in the current instance or resource space	≥ 0	Count
	iotda_south_messageUp_successCount	Successful MQTT Message Reporting Times	Number of times MQTT messages are successfully reported in the current instance or resource space	≥ 0	Count
	iotda_south_messageUp_failedCount	Failed MQTT Message Reporting Times	Number of times MQTT messages fail to be reported in the current instance or resource space	≥ 0	Count
AMQP transfers	iotda_amqp_forwarding_totalCount	AMQP Transfers	Number of AMQP transfers in the current instance or resource space	≥ 0	Count
	iotda_amqp_forwarding_successCount	Successful AMQP Transfers	Number of successful AMQP transfers in the current instance or resource space	≥ 0	Count
	iotda_amqp_forwarding_failedCount	Failed AMQP Transfers	Number of failed AMQP transfers in the current instance or resource space	≥ 0	Count
FunctionGraph transfers	iotda_functionGraph_forwarding_totalCount	FunctionGraph Transfers	Number of FunctionGraph transfers in the current instance or resource space	≥ 0	Count
	iotda_functionGraph_forwarding_successCount	Successful FunctionGraph Transfers	Number of successful FunctionGraph transfers in the current instance or resource space	≥ 0	Count

Category	Metric	Metric Name	Description	Value Range	Unit
	iotda_functionGraph_forwarding_failedCount	Failed FunctionGraph Transfers	Number of failed FunctionGraph transfers in the current instance or resource space	≥ 0	Count
MRS Kafka transfers	iotda_mrsKafka_forwarding_totalCount	MRS Kafka Transfers	Number of MRS Kafka transfers in the current instance or resource space	≥ 0	Count
	iotda_mrsKafka_forwarding_successCount	Successful MRS Kafka Transfers	Number of successful MRS Kafka transfers in the current instance or resource space	≥ 0	Count
	iotda_mrsKafka_forwarding_failedCount	Failed MRS Kafka Transfers	Number of failed MRS Kafka transfers in the current instance or resource space	≥ 0	Count
MQTT transfers	iotda_mqtt_forwarding_totalCount	MQTT Transfers	Number of MQTT transfers in the current instance or resource space	≥ 0	Count
	iotda_mqtt_forwarding_successCount	Successful MQTT Transfers	Number of successful MQTT transfers in the current instance or resource space	≥ 0	Count
	iotda_mqtt_forwarding_failedCount	Failed MQTT Transfers	Number of failed MQTT transfers in the current instance or resource space	≥ 0	Count
MySQL transfers	iotda_mysql_forwarding_totalCount	MySQL Transfers	Number of MySQL transfers in the current instance or resource space	≥ 0	Count
	iotda_mysql_forwarding_successCount	Successful MySQL Transfers	Number of successful MySQL transfers in the current instance or resource space	≥ 0	Count

Category	Metric	Metric Name	Description	Value Range	Unit
	iotda_mysql_forwarding_failedCount	Failed MySQL Transfers	Number of failed MySQL transfers in the current instance or resource space	≥ 0	Count
InfluxDB transfers	iotda_influxDB_forwarding_totalCount	InfluxDB Transfers	Number of InfluxDB transfers in the current instance or resource space	≥ 0	Count
	iotda_influxDB_forwarding_successCount	Successful InfluxDB Transfers	Number of successful InfluxDB transfers in the current instance or resource space	≥ 0	Count
	iotda_influxDB_forwarding_failedCount	Failed InfluxDB Transfers	Number of failed InfluxDB transfers in the current instance or resource space	≥ 0	Count
HTTP message pushes	iotda_http_forwarding_totalCount	HTTP Message Pushes	Number of HTTP message pushes in the current instance or resource space	≥ 0	Count
	iotda_http_forwarding_successCount	Successful HTTP Message Pushes	Number of times HTTP messages are successfully pushed in the current instance or resource space	≥ 0	Count
	iotda_http_forwarding_failedCount	Failed HTTP Message Pushes	Number of times HTTP messages failed to be pushed in the current instance or resource space	≥ 0	Count
OBS transfers	iotda_obs_forwarding_totalCount	OBS Transfers	Number of OBS transfers in the current instance or resource space	≥ 0	Count
	iotda_obs_forwarding_successCount	Successful OBS Transfers	Number of successful OBS transfers in the current instance or resource space	≥ 0	Count

Category	Metric	Metric Name	Description	Value Range	Unit
	iotda_obs_forwarding_failedCount	Failed OBS Transfers	Number of failed OBS transfers in the current instance or resource space	≥ 0	Count
DMS Kafka transfers	iotda_dmsKafka_forwarding_totalCount	DMS Kafka Transfers	Number of DMS Kafka transfers in the current instance or resource space	≥ 0	Count
	iotda_dmsKafka_forwarding_successCount	Successful DMS Kafka Transfers	Number of successful DMS Kafka transfers in the current instance or resource space	≥ 0	Count
	iotda_dmsKafka_forwarding_failedCount	Failed DMS Kafka Transfers	Number of failed DMS Kafka transfers in the current instance or resource space	≥ 0	Count
DIS transfers	iotda_dis_forwarding_totalCount	DIS Transfers	Number of DIS transfers in the current instance or resource space	≥ 0	Count
	iotda_dis_forwarding_successCount	Successful DIS Transfers	Number of successful DIS transfers in the current instance or resource space	≥ 0	Count
	iotda_dis_forwarding_failedCount	Failed DIS Transfers	Number of failed DIS transfers in the current instance or resource space	≥ 0	Count
ROMA Connect transfers	iotda_roma_forwarding_totalCount	ROMA Connect Transfers	Number of ROMA Connect transfers in the current instance or resource space	≥ 0	Count
	iotda_roma_forwarding_successCount	Successful ROMA Connect Transfers	Number of successful ROMA Connect transfers in the current instance or resource space	≥ 0	Count

Category	Metric	Metric Name	Description	Value Range	Unit
	iotda_roma_forwarding_failedCount	Failed ROMA Connect Transfers	Number of failed ROMA Connect transfers in the current instance or resource space	≥ 0	Count
LTS transfers	iotda_lts_forwarding_totalCount	LTS Transfers	Number of LTS transfers in the current instance or resource space	≥ 0	Count
	iotda_lts_forwarding_successCount	Successful LTS Transfers	Number of successful LTS transfers in the current instance or resource space	≥ 0	Count
	iotda_lts_forwarding_failedCount	Failed LTS Transfers	Number of failed LTS transfers in the current instance or resource space	≥ 0	Count
Device shadow querying	iotda_query_shadow_totalCount	Total Device Shadow Queries	Total number of device shadow queries in the current instance or resource space	≥ 0	Count
	iotda_query_shadow_successCount	Successful Device Shadow Queries	Number of successful device shadow queries in the current instance or resource space	≥ 0	Count
	iotda_query_shadow_failedCount	Failed Device Shadow Queries	Number of failed device shadow queries in the current instance or resource space	≥ 0	Count
Message delivery	iotda_message_down_totalCount	Total Messages Delivered	Total number of messages delivered in the current instance or resource space	≥ 0	Count
	iotda_message_down_successCount	Messages Successfully Delivered	Number of messages that are successfully delivered in the current instance or resource space	≥ 0	Count

Category	Metric	Metric Name	Description	Value Range	Unit
	iotda_message_down_failedCount	Messages Failing to Be Delivered	Number of messages that fail to be delivered in the current instance or resource space	≥ 0	Count
Property setting	iotda_properties_set_totalCount	Total Property Settings	Total number of property settings in the current instance or resource space	≥ 0	Count
	iotda_properties_set_successCount	Successful Property Settings	Number of successful property settings in the current instance or resource space	≥ 0	Count
	iotda_properties_set_failedCount	Failed Property Settings	Number of failed property settings in the current instance or resource space	≥ 0	Count
Property queries	iotda_properties_query_totalCount	Total Property Queries	Total number of property queries in the current instance or resource space	≥ 0	Count
	iotda_properties_query_successCount	Successful Property Queries	Number of successful property queries in the current instance or resource space	≥ 0	Count
	iotda_properties_query_failedCount	Failed Property Queries	Number of failed property queries in the current instance or resource space	≥ 0	Count
Commands	iotda_command_totalCount	Total Commands	Total number of commands in the current instance or resource space	≥ 0	Count
	iotda_command_successCount	Successful Command Executions	Number of commands that are successfully executed in the current instance or resource space	≥ 0	Count

Category	Metric	Metric Name	Description	Value Range	Unit
	iotda_command_failedCount	Failed Command Executions	Number of commands that fail to be executed in the current instance or resource space	≥ 0	Count
Flow control data	iotda_flowcontrol_north_failedCount	Flow Control Failures During API Invocation	Number of flow control failures during API invocation in the current instance	≥ 0	Count
	iotda_flowcontrol_south_failedCount	Flow Control Failures During Southbound Invocation	Number of flow control failures during data reporting in the current instance	≥ 0	Count
	iotda_flowcontrol_forwarding_failedCount	Flow Control Failures During Data Forwarding	Number of flow control failures during data forwarding in the current instance	≥ 0	Count
Concurrent online devices	iotda_max_online_devices_totalCount	Maximum Concurrent Online Devices per Day	Maximum number of concurrent online devices per day in the current instance	≥ 0	Count
	iotda_max_online_devices_percent	Percentage of Maximum Concurrent Online Devices	Percentage of the maximum number of concurrent online devices to the total number of devices per day in the current instance	0–100	%
TPS messages	iotda_mqtt_publishIn_totalCount	TPS Messages Reported Using MQTT	Number of TPS messages reported in the current instance	≥ 0	Count
	iotda_mqtt_connect_totalCount	MQTT Connections	Number of MQTT connections in the current instance	≥ 0	Count
	iotda_http_apigIn_totalCount	TPS Messages of Northbound API	Number of TPS messages of the northbound API in the current instance	≥ 0	Count

Category	Metric	Metric Name	Description	Value Range	Unit
Northbound API statistics	iotda_http_request_totalCount	Total Requests of Northbound API	Total number of requests of the northbound API in the current instance	≥ 0	Count
	iotda_http_request_failedCount	Failed Requests of Northbound API	Number of failed requests of the northbound API in the current instance	≥ 0	Count
	iotda_http_request_successCount	Successful Requests of Northbound API	Number of successful requests of the northbound API in the current instance	≥ 0	Count
MongoDB transfers	iotda_mongodb_forwarding_totalCount	Total MongoDB Transfers	Total number of MongoDB transfers in the current instance	≥ 0	Count
	iotda_mongodb_forwarding_successCount	Successful MongoDB Transfers	Number of successful MongoDB transfers in the current instance	≥ 0	Count
	iotda_mongodb_forwarding_failedCount	Failed MongoDB Transfers	Number of failed MongoDB transfers in the current instance	≥ 0	Count
bcs-fabric transfers	iotda_bcs_fabric_forwarding_totalCount	Total bcs-fabric transfers	Total number of bcs-fabric transfers in the current instance	≥ 0	Count
	iotda_bcs_fabric_forwarding_successCount	Successful bcs-fabric Transfers	Number of successful bcs-fabric transfers in the current instance	≥ 0	Count
	iotda_bcs_fabric_forwarding_failedCount	Failed bcs-fabric Transfers	Number of failed bcs-fabric transfers in the current instance	≥ 0	Count
bcs-huawei transfers	iotda_bcs_hw_forwarding_totalCount	Total bcs-huawei Transfers	Total number of bcs-huawei transfers in the current instance	≥ 0	Count

Category	Metric	Metric Name	Description	Value Range	Unit
	iotda_bcs_hw_forwarding_successCount	Successful bcs-huawei Transfers	Number of successful bcs-huawei transfers in the current instance	≥ 0	Count
	iotda_bcs_hw_forwarding_failedCount	Failed bcs-huawei Transfers	Number of failed bcs-huawei transfers in the current instance	≥ 0	Count
AMQP messages	iotda_amqp_forwarding_backlog_message_count	Stacked AMQP Messages	Number of stacked AMQP messages in the current instance	≥ 0	Count
	iotda_amqp_forwarding_consume_rate	Stacked AMQP Messages per Second	Number of stacked AMQP messages per second in the current instance	≥ 0	Count
Messages	iotda_userMessage_count	Total Messages	Total number of messages in the current instance	≥ 0	Count
Online devices	max_online_device	Concurrent Online Devices	Number of concurrent online devices in the current instance	≥ 0	Count
Batch tasks	iotda_batchtask_success_count	Successful Batch Tasks	Number of successful batch tasks in the current instance	≥ 0	Count
	iotda_batchtask_failure_count	Failed Batch Tasks	Number of failed batch tasks in the current instance	≥ 0	Count
Total metrics	iotda_indicator	Total Metrics	Total number of metrics of the current instance	≥ 0	Count
Data package size	iotda_http_request_packageSize	Data Package Size	Data package size of the current instance	≥ 0	Count

Category	Metric	Metric Name	Description	Value Range	Unit
	iotda_south_packageSize	Southbound Data Package Size	Size of southbound data package of the current instance	≥ 0	Count

10.8 Basic Metrics: Node Exporter Metrics

This section describes the types, names, and meanings of metrics reported by Node Exporter to AOM.

Table 10-8 Node Exporter metrics

Job	Metric	Description
node-exporter	node_filesystem_size_bytes	Consumed space of a file system
	node_filesystem_readonly	Read-only file system
	node_filesystem_free_bytes	Remaining space of a file system
	node_filesystem_avail_bytes	File system space that is available for use
	node_cpu_seconds_total	Seconds each CPU spent doing each type of work
	node_network_receive_bytes_total	Total amount of received data
	node_network_receive_errs_total	Cumulative number of errors encountered during reception
	node_network_transmit_bytes_total	Total amount of transmitted data
	node_network_receive_packets_total	Cumulative number of packets received
	node_network_transmit_drop_total	Cumulative number of dropped packets during transmission
	node_network_transmit_errs_total	Cumulative number of errors encountered during transmission
	node_network_up	NIC status

Job	Metric	Description
	node_network_transmit_packets_total	Cumulative number of packets transmitted
	node_network_receive_drop_total	Cumulative number of packets dropped during reception
	go_gc_duration_seconds	This value is obtained by calling the **debug.ReadGCStats()** function. When this function is called, the **PauseQuantile** field of the **GCStats** structure is set to **5**. In this way, the function returns 5 GC pause time percentiles (the minimum percentile, 25%, 50%, 75%, and maximum percentile). Then, the Prometheus Go client creates a summary metric based on the returned GC pause time percentile, **NumGC**, and **PauseTotal**.
	node_load5	5-minute average CPU load
	node_filefd_allocated	Allocated file descriptors
	node_exporter_build_info	Node Exporter build information
	node_disk_written_bytes_total	Number of bytes that are written
	node_disk_writes_completed_total	Number of writes completed
	node_disk_write_time_seconds_total	Number of seconds spent by all writes
	node_nf_conntrack_entries	Number of currently allocated flow entries for connection tracking
	node_nf_conntrack_entries_limit	Maximum size of a connection tracking table
	node_processes_max_processes	PID limit
	node_processes_pids	Number of PIDs
	node_sockstat_TCP_alloc	Number of allocated TCP sockets
	node_sockstat_TCP_inuse	Number of TCP sockets in use
	node_sockstat_TCP_tw	Number of TCP sockets in the **TIME_WAIT** state

Job	Metric	Description
	node_timex_offset_seconds	Time offset
	node_timex_sync_status	Synchronization status of node clocks
	node_uname_info	Labeled system information as provided by the uname system call
	node_vmstat_pgfault	Number of page faults the system has made per second in **/proc/vmstat**
	node_vmstat_pgmajfault	Number of major faults per second in **/proc/vmstat**
	node_vmstat_pgpgin	Number of page in between main memory and block device in **/proc/vmstat**
	node_vmstat_pgpgout	Number of page out between main memory and block device in **/proc/vmstat**
	node_disk_reads_completed_total	Number of reads completed
	node_disk_read_time_seconds_total	Number of seconds spent by all reads
	process_cpu_seconds_total	The value is obtained based on the **utime** parameter (the number of ticks executed by the Go process in user mode) and the **stime** parameter (the number of ticks executed by the Go process in kernel mode, for example, during system invocation). Unit: jiffies, which measure the tick time between two system timer interruptions. process_cpu_seconds_total = (utime + stime)/USER_HZ. Based on the preceding formula, you can obtain the total time (unit: seconds) for a process to run on the OS.
	node_disk_read_bytes_total	Number of bytes that are read
	node_disk_io_time_weighted_seconds_total	The weighted number of seconds spent doing I/Os

Job	Metric	Description
	node_disk_io_time_seconds_total	Total seconds spent doing I/Os
	node_disk_io_now	Number of I/Os in progress
	node_context_switches_total	Number of context switches
	node_boot_time_seconds	Node boot time
	process_resident_memory_bytes	Resident set size (RSS), which is the memory actually used by a process. It includes the shared memory, but does not include the allocated but unused memory or swapped-out memory.
	node_intr_total	Number of interruptions that occurred
	node_load1	1-minute average CPU load
	go_goroutines	This value is obtained by calling **runtime.NumGoroutine()** and calculated based on the **sched** scheduler structure and global **allglen** variable. Fields in the **sched** structure may change concurrently. Therefore, the system checks whether the calculated value is less than **1**. If the value is less than **1**, the system returns **1**.
	scrape_duration_seconds	Total time spent in collecting metrics
	node_load15	15-minute average CPU load
	scrape_samples_post_metric_relabeling	Number of remaining samples after metrics are relabeled
	node_netstat_Tcp_PassiveOpens	Number of TCP connections that directly change from the **LISTEN** state to the **SYN-RCVD** state
	scrape_samples_scraped	Number of samples scraped
	node_netstat_Tcp_CurrEstab	Number of TCP connections in the **ESTABLISHED** or **CLOSE-WAIT** state
	scrape_series_added	Number of series added to the scrape target

Job	Metric	Description
	node_netstat_Tcp_ActiveOpens	Number of TCP connections that directly change from the **CLOSED** state to the **SYN-SENT** state
	node_memory_MemTotal_bytes	Total memory of a node
	node_memory_MemFree_bytes	Free memory of a node
	node_memory_MemAvailable_bytes	Available memory of a node
	node_memory_Cached_bytes	Memory for the node page cache
	up	Scrape target status
	node_memory_Buffers_bytes	Memory of the node buffer

10.9 Basic Metrics: Flink Metrics

This section describes the categories, names, and meanings of Flink metrics reported to AOM.

Table 10-9 Flink metrics

Category	Metric	Description	Unit
CPU	flink_jobmanager_Status_JVM_CPU_Load	CPU load of the JVM in JobManager	N/A
	flink_jobmanager_Status_JVM_CPU_Time	CPU time of the JVM in JobManager	N/A
	flink_jobmanager_Status_ProcessTree_CPU_Usage	CPU usage of the JVM in JobManager	N/A
	flink_taskmanager_Status_JVM_CPU_Load	CPU load of the JVM in TaskManager	N/A
	flink_taskmanager_Status_JVM_CPU_Time	CPU time of the JVM in TaskManager	N/A
	flink_taskmanager_Status_ProcessTree_CPU_Usage	CPU usage of the JVM in TaskManager	N/A
Memory	flink_jobmanager_Status_JVM_Memory_Heap_Used	Used heap memory of JobManager	Bytes

Category	Metric	Description	Unit
	flink_jobmanager_Status_JVM_Memory_Heap_Committed	Available JVM heap memory of JobManager	Bytes
	flink_jobmanager_Status_JVM_Memory_Heap_Max	Maximum heap memory that can be used for memory management in JobManager	Bytes
	flink_jobmanager_Status_JVM_Memory_NonHeap_Used	Used off-heap memory of JobManager	Bytes
	flink_jobmanager_Status_JVM_Memory_NonHeap_Committed	Available JVM off-heap memory of JobManager	Bytes
	flink_jobmanager_Status_JVM_Memory_NonHeap_Max	Maximum off-heap memory that can be used for memory management in JobManager	Bytes
	flink_jobmanager_Status_JVM_Memory_Metaspace_Used	Used memory of the JobManager metaspace memory pool	Bytes
	flink_jobmanager_Status_JVM_Memory_Metaspace_Committed	Available JVM memory of the JobManager metaspace memory pool	Bytes
	flink_jobmanager_Status_JVM_Memory_Metaspace_Max	Maximum memory that can be used in the JobManager metaspace memory pool	Bytes
	flink_jobmanager_Status_JVM_Memory_Direct_Count	Number of buffers in the direct buffer pool of JobManager	N/A
	flink_jobmanager_Status_JVM_Memory_Direct_MemoryUsed	Memory for the direct buffer pool in JobManager	Bytes
	flink_jobmanager_Status_JVM_Memory_Direct_TotalCapacity	Total capacity of all buffers in the direct buffer pool of JobManager	Bytes
	flink_jobmanager_Status_JVM_Memory_Mapped_Count	Number of buffers in the mapped buffer pool of JobManager	N/A

Category	Metric	Description	Unit
	flink_jobmanager_Status_JVM_Memory_Mapped_MemoryUsed	Memory for the mapped buffer pool in JobManager	Bytes
	flink_jobmanager_Status_JVM_Memory_Mapped_TotalCapacity	Total capacity of all buffers in the mapped buffer pool of JobManager	Bytes
	flink_jobmanager_Status_Flink_Memory_Managed_Used	Managed memory that has been used in JobManager	Bytes
	flink_jobmanager_Status_Flink_Memory_Managed_Total	Total managed memory of JobManager	Bytes
	flink_taskmanager_Status_JVM_Memory_Heap_Used	Used heap memory of TaskManager	Bytes
	flink_taskmanager_Status_JVM_Memory_Heap_Committed	Available JVM heap memory of TaskManager	Bytes
	flink_taskmanager_Status_JVM_Memory_Heap_Max	Maximum heap memory that can be used for memory management in TaskManager	Bytes
	flink_taskmanager_Status_JVM_Memory_NonHeap_Used	Used off-heap memory of TaskManager	Bytes
	flink_taskmanager_Status_JVM_Memory_NonHeap_Committed	Available JVM off-heap memory of TaskManager	Bytes
	flink_taskmanager_Status_JVM_Memory_NonHeap_Max	Maximum off-heap memory that can be used for memory management in TaskManager	Bytes
	flink_taskmanager_Status_JVM_Memory_Metaspace_Used	Used memory of the TaskManager metaspace memory pool	Bytes
	flink_taskmanager_Status_JVM_Memory_Metaspace_Committed	Available JVM memory of the TaskManager metaspace memory pool	Bytes

Category	Metric	Description	Unit
	flink_taskmanager_Status_JVM_Memory_Metaspace_Max	Maximum memory that can be used in the TaskManager metaspace memory pool	Bytes
	flink_taskmanager_Status_JVM_Memory_Direct_Count	Number of buffers in the direct buffer pool of TaskManager	N/A
	flink_taskmanager_Status_JVM_Memory_Direct_MemoryUsed	Memory for the direct buffer pool in TaskManager	Bytes
	flink_taskmanager_Status_JVM_Memory_Direct_TotalCapacity	Total capacity of all buffers in the direct buffer pool of TaskManager	Bytes
	flink_taskmanager_Status_JVM_Memory_Mapped_Count	Number of buffers in the mapped buffer pool of TaskManager	N/A
	flink_taskmanager_Status_JVM_Memory_Mapped_MemoryUsed	Memory for the mapped buffer pool in TaskManager	Bytes
	flink_taskmanager_Status_JVM_Memory_Mapped_TotalCapacity	Total capacity of all buffers in the mapped buffer pool of TaskManager	Bytes
	flink_taskmanager_Status_Flink_Memory_Managed_Used	Managed memory that has been used in TaskManager	Bytes
	flink_taskmanager_Status_Flink_Memory_Managed_Total	Total managed memory of TaskManager	Bytes
	flink_taskmanager_Status_ProcessTree_Memory_RSS	Memory of the whole process in the Linux system	Bytes
Threads	flink_jobmanager_Status_JVM_Threads_Count	Total number of active threads in JobManager	Count
	flink_taskmanager_Status_JVM_Threads_Count	Total number of active threads in TaskManager	Count

Category	Metric	Description	Unit
Garbage collection	flink_jobmanager_Status_JVM_GarbageCollector_ConcurrentMarkSweep_Count	Number of garbage collection (GC) times of the JobManager Concurrent Mark Sweep (CMS) collector	Count
	flink_jobmanager_Status_JVM_GarbageCollector_ConcurrentMarkSweep_Time	Total time required for the JobManager CMS collector to collect garbage	ms
	flink_jobmanager_Status_JVM_GarbageCollector_ParNew_Count	Number of JobManager GC times	Count
	flink_jobmanager_Status_JVM_GarbageCollector_ParNew_Time	Each GC duration of JobManager	ms
	flink_taskmanager_Status_JVM_GarbageCollector_ConcurrentMarkSweep_Count	Number of GC times of the TaskManager CMS collector	Count
	flink_taskmanager_Status_JVM_GarbageCollector_ConcurrentMarkSweep_Time	Total time required for the TaskManager CMS collector to collect garbage	ms
	flink_taskmanager_Status_JVM_GarbageCollector_ParNew_Count	Number of TaskManager GC times	Count
	flink_taskmanager_Status_JVM_GarbageCollector_ParNew_Time	Each GC duration of TaskManager	ms
Class loader	flink_jobmanager_Status_JVM_ClassLoader_ClassesLoaded	Total number of classes that JobManager has loaded since the JVM started	N/A
	flink_jobmanager_Status_JVM_ClassLoader_ClassesUnloaded	Total number of classes that JobManager has unloaded since the JVM started	N/A
	flink_taskmanager_Status_JVM_ClassLoader_ClassesLoaded	Total number of classes that TaskManager has loaded since the JVM started	N/A

Category	Metric	Description	Unit
	flink_taskmanager_Status_JVM_ClassLoader_ClassesUnloaded	Total number of classes that TaskManager has unloaded since the JVM started	N/A
Network	flink_taskmanager_Status_Network_AvailableMemorySegments	Number of unused memory segments of TaskManager	N/A
	flink_taskmanager_Status_Network_TotalMemorySegments	Total number of allocated memory segments of TaskManager	N/A
Default shuffle service	flink_taskmanager_Status_Shuffle_Netty_AvailableMemorySegments	Number of unused memory segments of TaskManager	N/A
	flink_taskmanager_Status_Shuffle_Netty_UsedMemorySegments	Number of used memory segments of TaskManager	N/A
	flink_taskmanager_Status_Shuffle_Netty_TotalMemorySegments	Number of allocated memory segments of TaskManager	N/A
	flink_taskmanager_Status_Shuffle_Netty_AvailableMemory	Unused memory of TaskManager	Bytes
	flink_taskmanager_Status_Shuffle_Netty_UsedMemory	Used memory of TaskManager	Bytes
	flink_taskmanager_Status_Shuffle_Netty_TotalMemory	Allocated memory of TaskManager	Bytes
Availability	flink_jobmanager_job_numRestarts	Total number of restarts since job submission	Count
Checkpoint	flink_jobmanager_job_lastCheckpointDuration	Time taken to complete the latest checkpoint	ms
	flink_jobmanager_job_lastCheckpointSize	Size of the latest checkpoint. If incremental checkpoints are enabled or logs are changed, this metric may be different from **lastCheckpointFullSize**.	Bytes

Category	Metric	Description	Unit
	flink_jobmanager_job_numberOfInProgressCheck-points	Number of checkpoints that are in progress	Count
	flink_jobmanager_job_numberOfCompletedCheck-points	Number of checkpoints that are completed	Count
	flink_jobmanager_job_numberOfFailedCheckpoints	Number of failed checkpoints	Count
	flink_jobmanager_job_totalNumberOfCheckpoints	Total number of checkpoints	Count
I/O	flink_taskmanager_job_task_numBytesOut	Total number of bytes output by a task	Bytes
	flink_taskmanager_job_task_numBytesOutPerSecond	Total number of bytes output by a task per second	Bytes/s
	flink_taskmanager_job_task_isBackPressured	Whether a backpressure event occurs	N/A
	flink_taskmanager_job_task_numRecordsIn	Total number of records received by a task	Count
	flink_taskmanager_job_task_numRecordsInPerSecond	Total number of records received by a task per second	Records/s
	flink_taskmanager_job_task_numBytesIn	Number of bytes received by a task	Bytes
	flink_taskmanager_job_task_numBytesInPerSecond	Number of bytes received by a task per second	Bytes/s
	flink_taskmanager_job_task_numRecordsOut	Total number of records sent by a task	Count
	flink_taskmanager_job_task_numRecordsOutPerSecond	Total number of records sent by a task per second	Records/s
	flink_taskmanager_job_task_operator_numRecordsIn	Total number of records received by an operator	Count
	flink_taskmanager_job_task_operator_numRecordsInPerSecond	Total number of records received by an operator per second	Records/s
	flink_taskmanager_job_task_operator_numRecordsOut	Total number of records sent by an operator	Count

Category	Metric	Description	Unit
	flink_taskmanager_job_task_operator_numRecordsOutPerSecond	Total number of records sent by an operator per second	Records/s
	flink_taskmanager_job_task_operator_source_numRecordsIn	Total number of records input to the source	Count
	flink_taskmanager_job_task_operator_sink_numRecordsOut	Total number of records output from the sink	Count
	flink_taskmanager_job_task_operator_source_numRecordsInPerSecond	Number of records input to the source per second	Records/s
	flink_taskmanager_job_task_operator_sink_numRecordsOutPerSecond	Number of records output from the sink per second	Records/s
Kafka connector	flink_taskmanager_job_task_operator_currentEmitEventTimeLag	Interval between the data event time and the time when the data leaves the source	ms
	flink_taskmanager_job_task_operator_currentFetchEventTimeLag	Interval between the data event time and the time when the data enters the source	ms
	flink_taskmanager_job_task_operator_sourceIdleTime	Source idle duration	ms
	flink_taskmanager_job_task_operator_pendingRecords	Number of data records that have not been pulled by the source	Count

10.10 Basic Metrics: CCI 2.0 Metrics

This section describes the categories, names, and meanings of CCI 2.0 metrics reported to AOM.

Table 10-10 CCI 2.0 metrics

Category	Metric Name	Description
CPU	container_cpu_system_seconds_total	Cumulative container CPU system time

Category	Metric Name	Description
	container_cpu_usage_seconds_total	Cumulative CPU time consumed by a container in core-seconds
	container_cpu_user_seconds_total	Cumulative user CPU time
	container_cpu_cfs_periods_total	Number of elapsed enforcement period intervals
	container_cpu_cfs_throttled_periods_total	Number of throttled period intervals
	container_cpu_cfs_throttled_seconds_total	Total duration a container has been throttled
File system or disk	container_fs_inodes_free	Number of available inodes in a file system
	container_fs_usage_bytes	File system usage
	container_fs_inodes_total	Number of inodes in a file system
	container_fs_io_current	Number of I/Os currently in progress in a disk or file system
	container_fs_io_time_seconds_total	Cumulative seconds spent on doing I/Os by a disk or file system
	container_fs_io_time_weighted_seconds_total	Cumulative weighted I/O time of a disk or file system
	container_fs_limit_bytes	Total disk or file system capacity that can be consumed by a container
	container_fs_reads_bytes_total	Cumulative amount of disk or file system data read by a container
	container_fs_read_seconds_total	Cumulative number of seconds a container spent on reading disk or file system data
	container_fs_reads_merged_total	Cumulative number of merged disk or file system reads made by a container
	container_fs_reads_total	Cumulative number of disk or file system reads completed by a container

Category	Metric Name	Description
	container_fs_sector_reads_total	Cumulative number of disk or file system sector reads completed by a container
	container_fs_sector_writes_total	Cumulative number of disk or file system sector writes completed by a container
	container_fs_writes_bytes_total	Total amount of data written by a container to a disk or file system
	container_fs_write_seconds_total	Cumulative number of seconds a container spent on writing data to a disk or file system
	container_fs_writes_merged_total	Cumulative number of merged container writes to a disk or file system
	container_fs_writes_total	Cumulative number of disk or file system writes completed by a container
	container_blkio_device_usage_total	Blkio device bytes usage
Memory	container_memory_failures_total	Cumulative number of container memory allocation failures
	container_memory_failcnt	Number of memory usage hits limits
	container_memory_cache	Memory used for the page cache of a container
	container_memory_mapped_file	Size of a container memory mapped file
	container_memory_max_usage_bytes	Maximum memory usage recorded for a container
	container_memory_rss	Size of the resident memory set for a container
	container_memory_swap	Container swap usage
	container_memory_usage_bytes	Current memory usage of a container
	container_memory_working_set_bytes	Memory usage of the working set of a container

Category	Metric Name	Description
Network	container_network_receive_bytes_total	Total volume of data received by a container network
	container_network_receive_errors_total	Cumulative number of errors encountered during reception
	container_network_receive_packets_dropped_total	Cumulative number of packets dropped during reception
	container_network_receive_packets_total	Cumulative number of packets received
	container_network_transmit_bytes_total	Total volume of data transmitted on a container network
	container_network_transmit_errors_total	Cumulative number of errors encountered during transmission
	container_network_transmit_packets_dropped_total	Cumulative number of packets dropped during transmission
	container_network_transmit_packets_total	Cumulative number of packets transmitted
Container spec or status	container_processes	Number of processes running inside a container
	container_sockets	Number of open sockets for a container
	container_file_descriptors	Number of open file descriptors for a container
	container_threads	Number of threads running inside a container
	container_threads_max	Maximum number of threads allowed inside a container
	container_ulimits_soft	Soft ulimit value of process 1 in a container. Unlimited if the value is **-1**, except priority and nice.
	container_spec_cpu_period	CPU period of a container
	container_spec_cpu_shares	CPU share of a container
	container_spec_cpu_quota	CPU quota of a container
	container_spec_memory_limit_bytes	Memory limit for a container
	container_spec_memory_reservation_limit_bytes	Memory reservation limit for a container

Category	Metric Name	Description
	container_spec_memory_swap_limit_bytes	Memory swap limit for a container
	container_start_time_seconds	Running time of a container
	container_last_seen	Last time a container was seen by an exporter
Object status metrics	kube_pod_info	Pod information
	kube_pod_owner	Information about the pod's owner
	kube_pod_container_resource_limits	Number of container resource limits
	kube_replicaset_owner	Information about the ReplicaSet's owner

10.11 Basic Metrics: CAE Metrics

This section describes the types, names, and meanings of Cloud Application Engine (CAE) metrics reported to AOM.

Table 10-11 CAE metrics

Category	Metric Name	Description
job:kubernetes-cadvisor	cadvisor_version_info	Version information of the cAdvisor module (a Docker container monitoring tool) in Prometheus
	container_cpu_cfs_periods_total	Number of elapsed enforcement period intervals
	container_cpu_cfs_throttled_periods_total	Number of throttled period intervals
	container_cpu_cfs_throttled_seconds_total	Total duration a container has been throttled
	container_cpu_load_average_10s	Value of container CPU load average over the last 10 seconds
	machine_cpu_physical_cores	Number of physical CPU cores
	machine_cpu_sockets	Number of CPU slots on a host

Category	Metric Name	Description
	machine_nvm_avg_power_budget_watts	Average power budget for non-volatile memory (NVM)
	machine_nvm_capacity	NVM of a host (usually a server or node)
	machine_scrape_error	Error encountered when collecting metrics from the target host (usually a server or container)
	container_cpu_system_seconds_total	Cumulative container CPU system time
	container_cpu_usage_seconds_total	Cumulative CPU time consumed by a container in core-seconds
	container_cpu_user_seconds_total	Cumulative user CPU time
	container_file_descriptors	Number of open file descriptors for a container
	container_fs_inodes_free	Number of available inodes in a file system
	container_fs_inodes_total	Number of inodes in a file system
	container_fs_io_current	Number of I/Os currently in progress in a disk or file system
	container_fs_io_time_seconds_total	Cumulative seconds spent on doing I/Os by a disk or file system
	container_fs_io_time_weighted_seconds_total	Cumulative weighted I/O time of a disk or file system
	container_fs_limit_bytes	Total disk or file system capacity that can be consumed by a container
	container_fs_read_seconds_total	Cumulative number of seconds a container spent on reading disk or file system data
	container_fs_reads_bytes_total	Cumulative amount of disk or file system data read by a container

Category	Metric Name	Description
	container_fs_reads_merged_total	Cumulative number of merged disk or file system reads made by a container
	container_fs_reads_total	Cumulative number of disk or file system reads completed by a container
	container_fs_sector_reads_total	Cumulative number of disk or file system sector reads completed by a container
	container_fs_sector_writes_total	Cumulative number of disk or file system sector writes completed by a container
	container_fs_usage_bytes	File system usage
	container_fs_write_seconds_total	Cumulative number of seconds a container spent on writing data to a disk or file system
	container_fs_writes_bytes_total	Total amount of data written by a container to a disk or file system
	container_fs_writes_merged_total	Cumulative number of merged container writes to a disk or file system
	container_fs_writes_total	Cumulative number of disk or file system writes completed by a container
	container_last_seen	Last time a container was seen by an exporter
	container_memory_cache	Memory used for the page cache of a container
	container_memory_failcnt	Number of memory usage hits limits
	container_memory_mapped_file	Size of a container memory mapped file
	container_memory_max_usage_bytes	Maximum memory usage recorded for a container
	container_memory_rss	Size of the resident memory set for a container
	container_memory_swap	Container swap usage

Category	Metric Name	Description
	container_memory_usage_bytes	Current memory usage of a container
	container_memory_working_set_bytes	Memory usage of the working set of a container
	container_network_receive_bytes_total	Total volume of data received by a container network
	container_network_receive_errors_total	Cumulative number of errors encountered during reception
	container_network_receive_packets_dropped_total	Cumulative number of packets dropped during reception
	container_network_receive_packets_total	Cumulative number of packets received
	container_network_transmit_bytes_total	Total volume of data transmitted on a container network
	container_network_transmit_errors_total	Cumulative number of errors encountered during transmission
	container_network_transmit_packets_dropped_total	Cumulative number of packets dropped during transmission
	container_network_transmit_packets_total	Cumulative number of packets transmitted
	container_processes	Number of processes running inside a container
	container_scrape_error	Number of errors encountered when collecting metrics from a container
	container_sockets	Number of open sockets for a container
	container_spec_cpu_period	CPU period of a container
	container_spec_cpu_quota	CPU quota of a container
	container_spec_cpu_shares	CPU share of a container
	container_spec_memory_limit_bytes	Memory limit for a container

Category	Metric Name	Description
	container_spec_memory_reservation_limit_bytes	Memory reservation limit for a container
	container_spec_memory_swap_limit_bytes	Memory swap limit for a container
	container_start_time_seconds	Running time of a container
	container_tasks_state	Number of tasks in the specified state, such as sleeping, running, stopped, uninterruptible, or ioawaiting
	container_threads	Number of threads running inside a container
	container_threads_max	Maximum number of threads allowed inside a container
	container_ulimits_soft	Soft ulimit value of process 1 in a container. Unlimited if the value is **-1**, except priority and nice.
	machine_cpu_cores	Number of CPU cores of a physical machine or VM
	machine_memory_bytes	Total memory size of a physical machine or VM
	scrape_duration_seconds	Time taken to obtain monitoring data (that is, "scrape" data) from the target resource
	scrape_samples_post_metric_relabeling	Number of samples collected after relabeling
	scrape_samples_scraped	Number of samples scraped in a specific period
	scrape_series_added	Performance of the monitoring data collection system (such as Prometheus)

Category	Metric Name	Description
	up	Running time of a system component or service. For example, in a containerized environment (such as a Docker container), **up** indicates the startup time of a container.
job:kubernetes-node-exporter	node_filesystem_size_bytes	Consumed space of a file system
	node_memory_Buffers_bytes	Memory of the node buffer
	node_cpu_seconds_total	Seconds each CPU spent doing each type of work
	node_filesystem_avail_bytes	File system space that is available for use
	node_memory_Cached_bytes	Memory for the node page cache
	node_memory_MemFree_bytes	Free memory of a node
	node_memory_MemTotal_bytes	Total memory of a node
	node_network_receive_bytes_total	Total amount of received data
	node_network_transmit_bytes_total	Total amount of transmitted data
	node_disk_reads_completed_total	Number of disk reads completed
	node_disk_writes_completed_total	Number of disk writes completed
Aggregated metrics	cae_container_cpu_usage	CPU usage of a container
	cae_container_downlink_bps	Downlink bandwidth of a container (that is, the rate for transmitting data from the container to the external network)
	cae_container_fs_reads_bytes_total	Total number of bytes read by a container from a file system
	cae_container_fs_writes_bytes_total	Number of bytes written by a container to a file system

Category	Metric Name	Description
	cae_container_memory_usage	Memory usage of a container
	cae_container_uplink_bps	Uplink bandwidth rate (bit/s) of a container
	cae_node_cup_usage	CPU usage of a node (usually a Kubernetes node or physical/virtual server)
	cae_node_disk_reads_bytes	Number of bytes read from the disk of a node (usually a host or server)
	cae_node_disk_usage	Disk usage of a node (usually a Kubernetes node or server)
	cae_node_disk_writes_bytes	Number of bytes written to the disk of a node (usually a host or server)
	cae_node_downlink_bps	Downlink bandwidth rate (bit/s) of a node (usually a server or VM)
	cae_node_memory_usage	Memory usage of a node (usually a Kubernetes node or physical/virtual server)
	cae_node_uplink_bps	Uplink bandwidth rate (bit/s) of a node in a Kubernetes cluster

10.12 Metric Dimensions

Dimensions of VM Metrics Reported by ICAgents

Table 10-12 Dimensions of VM metrics reported by ICAgents

Category	Metric Dimension	Description
Network metrics	clusterId	Cluster ID
	hostID	Host ID
	nameSpace	Cluster namespace
	netDevice	NIC name

Category	Metric Dimension	Description
	nodeIP	Host IP address
	nodeName	Host name
Disk metrics	clusterId	Cluster ID
	diskDevice	Disk name
	hostID	Host ID
	nameSpace	Cluster namespace
	nodeIP	Host IP address
	nodeName	Host name
Disk partition metrics	diskPartition	Partition disk
	diskPartitionType	Disk partition type
File system metrics	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
Host metrics	clusterId	Cluster ID
	clusterName	Cluster name
	gpuName	GPU name
	gpuID	GPU ID
	npuName	NPU name
	npuID	NPU ID
	hostID	Host ID
	nameSpace	Cluster namespace
	nodeIP	Host IP address
	hostName	Host name
Cluster metrics	clusterId	Cluster ID

Category	Metric Dimension	Description
	clusterName	Cluster name
	projectId	Project ID
Container metrics	appID	Service ID
	appName	Service name
	clusterId	Cluster ID
	clusterName	Cluster name
	containerID	Container ID
	containerName	Container name
	deploymentName	Workload name
	kind	Application type
	nameSpace	Cluster namespace
	podID	Instance ID
	podIP	Pod IP address
	podName	Instance name
	serviceID	Inventory ID
	nodename	Host name
	nodeIP	Host IP address
	virtualServiceName	Istio virtual service name
	gpuID	GPU ID
	npuName	NPU name
	npuID	NPU ID
Process metrics	appName	Service name
	clusterId	Cluster ID
	clusterName	Cluster name
	nameSpace	Cluster namespace
	processID	Process ID
	processName	Process name
	serviceID	Inventory ID

11 Security

11.1 Shared Responsibilities

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

Unlike traditional on-premises data centers, cloud computing separates operators from users. This approach not only enhances flexibility and control for users but also greatly reduces their operational workload. For this reason, cloud security cannot be fully ensured by one party. Cloud security requires joint efforts of Huawei Cloud and you, as shown in **Figure 11-1**.

- **Huawei Cloud**: Huawei Cloud is responsible for infrastructure security, including security and compliance, regardless of cloud service categories. The infrastructure consists of physical data centers, which house compute, storage, and network resources, virtualization platforms, and cloud services Huawei Cloud provides for you. In PaaS and SaaS scenarios, Huawei Cloud is responsible for security settings, vulnerability remediation, security controls, and detecting any intrusions into the network where your services or Huawei Cloud components are deployed.
- **Customer**: As our customer, your ownership of and control over your data assets will not be transferred under any cloud service category. Without your explicit authorization, Huawei Cloud will not use or monetize your data, but you are responsible for protecting your data and managing identities and access. This includes ensuring the legal compliance of your data on the cloud, using secure credentials (such as strong passwords and multi-factor authentication), and properly managing those credentials, as well as monitoring and managing content security, looking out for abnormal account behavior, and responding to it, when discovered, in a timely manner.

Figure 11-1 Huawei Cloud shared security responsibility model

Cloud security responsibilities are determined by control, visibility, and availability. When you migrate services to the cloud, assets, such as devices, hardware, software, media, VMs, OSs, and data, are controlled by both you and Huawei Cloud. This means that your responsibilities depend on the cloud services you select. As shown in **Figure 11-1**, customers can select different cloud service types (such as IaaS, PaaS, and SaaS) based on their service requirements. As control over components varies across different cloud service categories, the responsibilities are shared differently.

- In on-premises scenarios, customers have full control over assets such as hardware, software, and data, so tenants are responsible for the security of all components.
- In IaaS scenarios, customers have control over all components except the underlying infrastructure. So, customers are responsible for securing these components. This includes ensuring the legal compliance of the applications, maintaining development and design security, and managing vulnerability remediation, configuration security, and security controls for related components such as middleware, databases, and operating systems.
- In PaaS scenarios, customers are responsible for the applications they deploy, as well as the security settings and policies of the PaaS middleware, database, and network access under their control.
- In SaaS scenarios, customers have control over their content, accounts, and permissions. They need to protect their content, and properly configure and protect their accounts and permissions in compliance with laws and regulations.

On-premises (On-Prem): Software and IT infrastructure are deployed and managed by customers within their own data centers, rather than be deployed by remote cloud service providers.

Infrastructure as a Service (IaaS): Cloud service providers offer compute, network, storage, and more infrastructure services, including **Elastic Cloud Server (ECS)**, **Virtual Private Network (VPN)**, and **Object Storage Service (OBS)**.

Platform as a Service (PaaS): Cloud service providers deliver platforms required for application development and deployment, such as **ModelArts** and **GaussDB**. Customers do not need to maintain the underlying infrastructure.

Software as a Service (SaaS): Cloud service providers offer complete application software, such as **Huawei Cloud Meeting**. Customers use the software directly without the need to install the application, maintain it, or manage its underlying platform or infrastructure.

11.2 Access Control for AOM

Identity Authentication

Present your identity credential and undergo identity authentication no matter whether you access AOM through the console or by calling APIs. In addition, login protection and login authentication policies are provided to harden identity authentication security. Based on IAM, AOM supports three identity authentication modes: **Password Policy**, **Access Keys**, and **Temporary Access Key**. It also provides **Login Protection** and **Login Authentication Policy**.

Access Control

If you need to assign different permissions to employees in your enterprise to access your AOM resources, IAM is a good choice for fine-grained permissions management. IAM provides identity authentication, fine-grained permissions management, and access control. IAM helps you secure access to your Huawei Cloud resources. For details, see **13 Permissions Management**.

11.3 Data Protection

AOM takes different measures to keep data secure and reliable.

Table 11-1 AOM data protection methods and features

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

11.4 Audit and Logs

Audit

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

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

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

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

Figure 11-2 CTS

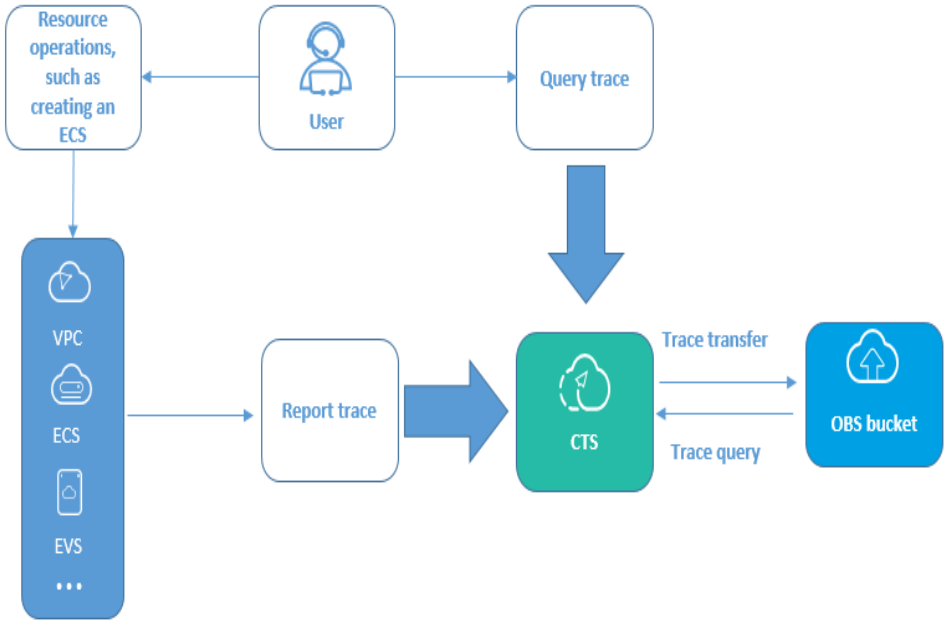

Logs

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

11.5 Resilience

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

Table 11-2 Reliability architecture of AOM

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

11.6 Security Risk Monitoring

AOM monitors security risks in various ways to ensure data security and reliability. For details, see **Table 11-3**.

Table 11-3 Monitoring security risks

Security Risk Monitoring	Description	Reference
Resource monitoring	AOM supports workload/cluster/host monitoring and metric browsing. It monitors your applications and cloud resources in real time and displays data, helping you quickly analyze application health.	• **Workload Monitoring** • **Cluster Monitoring** • **Host Monitoring** • **Metric Browsing**
Alarm management	AOM allows you to set alarm conditions for applications, resources, and services. When AOM or its external service is or may be abnormal, email, SMS, or WeCom notifications will be sent to specified personnel.	**Alarm Management**

11.7 Certificates

Compliance Certificates

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

Figure 11-3 Downloading compliance certificates

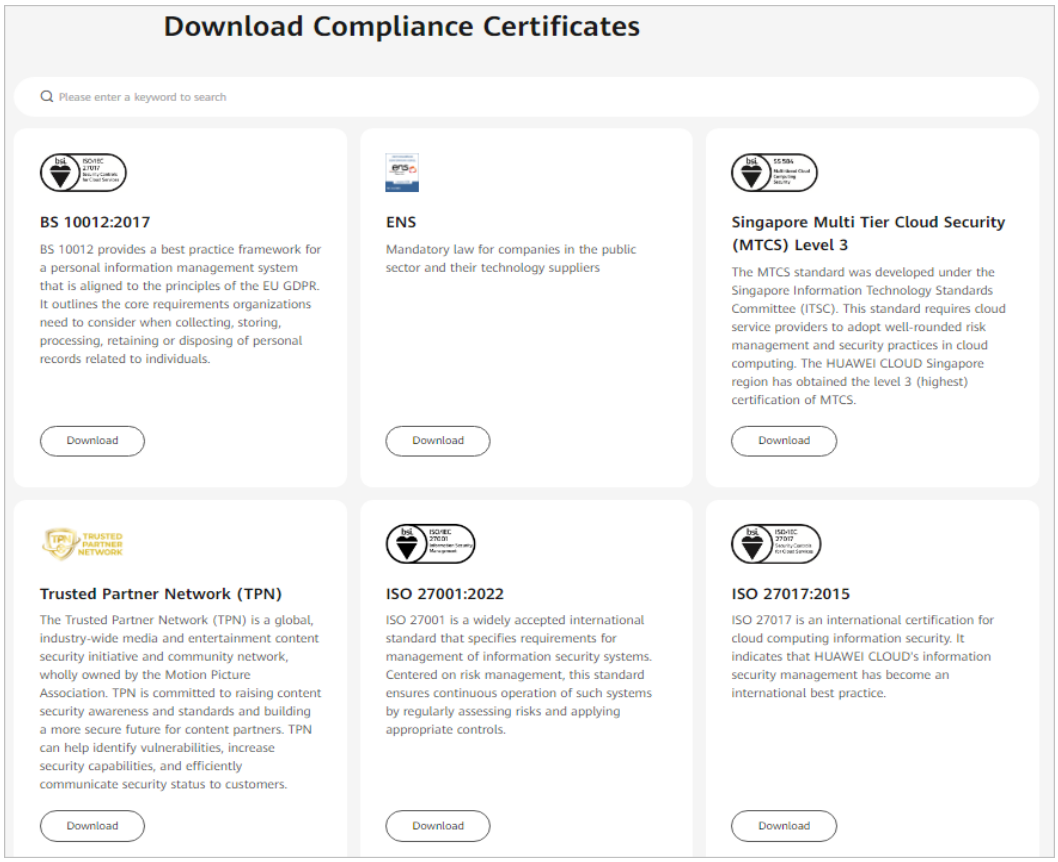

Resource Center

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

Figure 11-4 Resource center

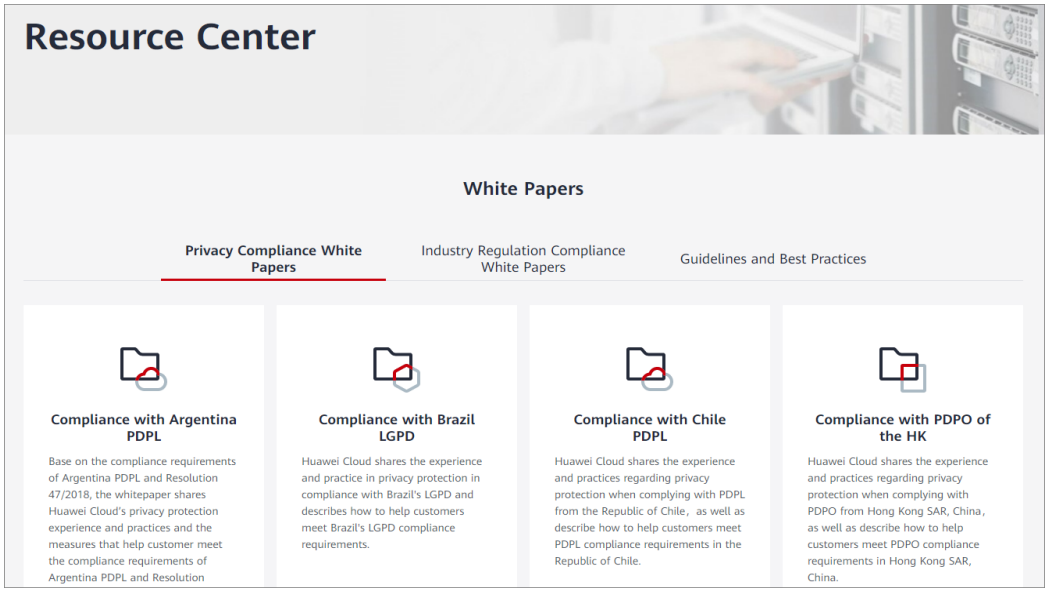

12 Basic Concepts

12.1 CMDB

Table 12-1 Basic concepts of CMDB

Terminology	Description
CMDB	CMDB manages relationships between applications and cloud resources and provides unified and timely resource environment management for AOM monitoring, automatic O&M, and APM.
Application topology	Hierarchical relationships between resources and applications. CMDB uses the "application + sub-application (optional) + component + environment" management model.
Enterprise project	An enterprise project can contain one or more applications.
Application	Basic object of CMDB and root node of the resource management model. An application represents a logical unit, which can be a project, product, or service. After an application is created, you can view the same application topology information in all regions. For example, an online shopping application includes user management, login verification, item list, and order management services.
Sub-application	Optional node for application management. An application can contain up to three layers of sub-applications. A sub-application can be regarded as a service, which is a group of components or microservices.

Terminology	Description
Component	Minimum unit of an application. It can be regarded as a program or middleware on which an application depends, such as Relational Database Service (RDS) and Distributed Message Service (DMS). Generally, a component is used together with environments. It can contain one or more environments. For example, an order application includes the function test environment, pressure test environment, pre-release environment, and live network environment.
Environment	Used to distinguish different environments or deployment regions of a component. Components or programs with different configurations form different environments. Each environment has its own region attribute. You can filter environments by region. You can also add one or more tags when creating an environment, and filter environments by tag. For example, a formal or test environment.
Environment tag	An attribute set for an environment. Multiple environments may have the same tag. You can filter required environments by tag. A tag can be added only to different environments of the same application.
Resource binding	You can bind a resource object to an environment of an application. A resource object instance of an application can belong to multiple environments.
Resource unbinding	If a component or environment changes and resources are not required, you can unbind them.
Resource transfer	If the component or environment to which a resource is bound changes, transfer the resource to the target node.

12.2 Resource Monitoring

Table 12-2 Basic concepts

Terminology	Description
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.

Terminology	Description
Host	Each host of AOM corresponds to a VM or physical machine. A host can be your own VM or physical machine, or an Elastic Cloud Service (ECS) or Bare Metal Server (BMS) purchased from Huawei Cloud. A host can be connected to AOM for monitoring only when its OS meets requirements and it is installed with an ICAgent.
Logs	AOM supports log collection, search, analysis, download, and dump. It also reports alarms based on keyword statistics and enables you to export reports, query SQL statements, and monitor data in real time.
Log traffic	Log traffic refers to the volume of logs reported per second. A maximum of 10 MB/s is supported for each tenant in a region. If the log traffic exceeds 10 MB/s, logs may be lost.
Alarms	Alarms are reported when AOM, ServiceStage, APM, or CCE is abnormal or may cause exceptions. Alarms will cause service exceptions and need to be handled.
Events	Events generally carry important information. They are reported when AOM, ServiceStage, APM, or CCE encounters some changes. Events do not necessarily cause service exceptions. Events do not need to be handled.
Alarm clearance	There are two alarm clearance modes: ● Automatic clearance: After a fault is rectified, AOM automatically clears the corresponding alarm. ● Manual clearance: After a fault is rectified, AOM does not automatically clear the corresponding alarm. Instead, you need to manually clear the alarm.
Alarm rules	Alarm rules are classified into metric alarm rules and event alarm rules. ● Metric alarm rules monitor the usage of resources (such as hosts and components) in the environment in real time. ● If there are many resource alarms but you do not want to receive notifications too often, set event alarm rules to quickly identify specific types of resource usage problems.
Alarm notification	There are two alarm notification modes: ● Direct alarm reporting: When setting alarm notification rules, specify alarm notification recipients so that they can take measures to rectify faults in a timely manner. Alarms can be sent through email, DingTalk, WeCom, voice calls, and SMS. ● Alarm noise reduction: Select a grouping rule to reduce alarm noise.

Terminology	Description
Alarm notification rule	An alarm notification rule defines the action to be taken after an alarm is generated. It includes where the message is sent and in what form. You can specify a message destination by setting **an SMN topic**.
Prometheus instances	Logical units used to manage Prometheus data collection, storage, and analysis.
Prometheus probes	Deployed in the Kubernetes clusters on the user or cloud product side. Prometheus probes automatically discover targets, collect metrics, and remotely write data to databases.
Exporters	Collect monitoring data and regulate the data provided for external systems using the Prometheus monitoring function. Currently, hundreds of official or third-party exporters are available. For details, see **Exporters**.
Jobs	Configuration set for a group of targets. Jobs specify the capture interval, access limit, and other behavior for a group of targets.

12.3 Automation

Table 12-3 Basic concepts of Automation

Terminology	Description
Script management	Shell, Python, BAT, and PowerShell scripts are supported and multiple versions of a single script can be managed.
Job management	You can orchestrate atomic operations about scripts and files to form job templates for specific O&M automation scenarios. For example, initializing a service environment.
Execution plans	You can select one or more steps from job templates to form execution plans. These plans are instantiated objects of job templates.
Cloud service scenarios	Atomic change scenarios of cloud services. For example, restarting an ECS.
Standardized O&M	You can assemble scripts, file management, execution plans, and cloud service scenarios to form operation processes. These processes standardize changes in specific O&M scenarios.

Terminology	Description
Scenarios	Jobs and standard O&M processes are released as services. You only need to enter simple information to complete specific O&M, without understanding underlying operation logic.
Parameter library	It consists of global parameters defined in jobs and standardized O&M. Both strings and hosts are supported. Parameters can be shared in different steps. In addition, all parameters in the entire process and their functions can be clearly viewed.
OS account	Used to execute scripts and manage files.
Tool market	By default, tool cards are displayed by scenario. You can bring these cards online or offline from the market as required and configure a review process to improve security.

12.4 Collection Management

Table 12-4 Basic concepts of collection management

Terminology	Description
UniAgent	UniAgent manages the life cycle of plug-ins centrally and deliver instructions for operations such as script delivery or execution. It does not collect O&M data; instead, different plug-ins do so. Install, upgrade, and uninstall these plug-ins as required. More plug-ins (such as Cloud Eye and Host Security Service (HSS)) are coming soon.
AK/SK	Access key. You can install ICAgents using tenant-level AK/SK for easy log collection.
ICAgent	ICAgents collect metrics, logs, and application performance data. For the hosts purchased on the ECS or BMS console, manually install ICAgents. For the hosts that are purchased through CCE, ICAgents are automatically installed.
Installation host	You can deliver UniAgent installation instructions to hosts in batches through an installation host on AOM. After setting an installation host, you can remotely install UniAgents on other hosts in the same VPC.
Proxy area/ Proxy	To enable network communication between multiple clouds, purchase and configure an ECS as a proxy and bind an EIP to it. AOM delivers deployment and control instructions to remote hosts and receives O&M data through the proxy. A proxy area contains multiple proxies for high availability.

13 Permissions Management

If you need to assign different permissions to employees in your enterprise to access your purchased AOM resources on Huawei Cloud, Identity and Access Management (IAM) is a good choice for fine-grained permissions management. IAM provides identity authentication, fine-grained permissions management, and access control. IAM helps you secure access to your Huawei Cloud resources. If your HUAWEI ID does not require IAM for permissions management, you can skip this section.

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

With IAM, you can control access to specific Huawei Cloud resources. For example, if you want some software developers in your enterprise to use AOM resources but do not want them to delete resources, you can grant permission to use AOM but not permission to delete its resources.

IAM supports role/policy-based authorization and identity policy-based authorization.

The following table describes the differences between the two authorization models.

Table 13-1 Differences between role/policy-based and identity policy-based authorization

Name	Core Relationship	Permissions	Authorization Method	Description
Role/Policy-based authorization	User-permission-authorization scope	• System-defined roles • System-defined policy • Custom policies	Assign roles or policies to principals.	To authorize a user, you need to add it to a user group first and then specify the scope of authorization. It provides a limited number of condition keys and cannot meet the requirements of fine-grained permissions control. This method is suitable for small- and medium-sized enterprises.
Identity policy-based authorization	User-policy	• System-defined identity policies • Custom identity policies	• Assign identity policies to principals. • Attach identity policies to principals.	You can grant permissions directly to a user. A variety of key conditions are available for more fine-grained permissions control. However, this model requires a certain level of expertise and is suitable for medium- and large-sized enterprises.

Assume that you want to grant IAM users permission to create ECSs in CN North-Beijing4 and OBS buckets in CN South-Guangzhou. With role/policy-based authorization, the administrator needs to create two custom policies and assign both to the IAM users. With identity policy-based authorization, the administrator only needs to create one custom identity policy and configure the condition key **g:RequestedRegion** for the policy, and then attach the policy to the users or grant the users the access permissions to the specified regions. Identity policy-based authorization is more flexible than role/policy-based authorization.

Policies/identity policies and actions in the two authorization models are not interoperable. You are advised to use the identity policy-based authorization model. For details about system-defined permissions, see **Role/Policy-based Authorization** and **Identity Policy-based Authorization**.

For more information about IAM, see **IAM Service Overview**.

Role/Policy-based Authorization

AOM supports role/policy-based authorization. New IAM users do not have any permissions assigned by default. You need to first add them to one or more groups and then attach policies or roles to these groups. These users then inherit permissions from the groups and perform specified operations on cloud services.

AOM is a project-level service deployed and accessed in specific physical regions. When you set **Scope** to **Region-specific projects** and select the specified projects (for example, **ap-southeast-2**) in the specified regions (for example, **AP-Bangkok**), the users only have permissions for resources in the selected projects. If you set **Scope** to **All resources**, the users have permissions for resources 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.

Table 13-2 lists all the system permissions supported by AOM. System-defined policies in role/policy-based authorization are not interoperable with those in identity policy-based authorization.

Table 13-2 System permissions supported by AOM

Policy Name	Description	Type	Dependency Permissions
AOM FullAccess	Administrator permissions for AOM 2.0. Users granted these permissions can operate and use AOM.	System-defined policy	CCE FullAccess, DMS ReadOnlyAccess, CCE Namespace-level Permissions, LTS FullAccess **For CCE namespaces, users or user groups must be granted the administrator (cluster-admin) or custom permissions. If custom permissions are granted, the get, list, and update permissions must be included and the resources of configmaps, prometheuses, servicemonitors, podmonitors, and namespaces must also be specified. For details, see Namespace Permissions (Kubernetes RBAC-based).**

Policy Name	Description	Type	Dependency Permissions
AOM ReadOnlyAccess	Read-only permissions for AOM 2.0. Users granted these permissions can only view AOM data.	System-defined policy	CCE ReadOnlyAccess, DMS ReadOnlyAccess, CCE Namespace-level Permissions, LTS ReadOnlyAccess **For CCE namespaces, users or user groups must be granted the administrator (cluster-admin) or custom permissions. If custom permissions are granted, the get and list permissions must be included and the resources of configmaps, prometheuses, servicemonitors, podmonitors, and namespaces must also be specified. For details, see Namespace Permissions (Kubernetes RBAC-based).**

Table 13-3 lists the common operations supported by each system-defined policy of resource monitoring. Select policies as required.

Table 13-3 Common operations supported by each system-defined policy

Operation	AOM FullAccess	AOM ReadOnlyAccess
Creating an alarm rule	√	x
Modifying an alarm rule	√	x
Deleting an alarm rule	√	x

Operation	AOM FullAccess	AOM ReadOnlyAccess
Creating an alarm template	√	x
Modifying an alarm template	√	x
Deleting an alarm template	√	x
Creating an alarm notification rule	√	x
Modifying an alarm notification rule	√	x
Deleting an alarm notification rule	√	x
Creating a message template	√	x
Modifying a message template	√	x
Deleting a message template	√	x
Creating a grouping rule	√	x
Modifying a grouping rule	√	x
Deleting a grouping rule	√	x
Creating a suppression rule	√	x
Modifying a suppression rule	√	x
Deleting a suppression rule	√	x
Creating a silence rule	√	x
Modifying a silence rule	√	x
Deleting a silence rule	√	x
Creating a dashboard	√	x
Modifying a dashboard	√	x
Deleting a dashboard	√	x
Creating a Prometheus instance	√	x

Operation	AOM FullAccess	AOM ReadOnlyAccess
Modifying a Prometheus instance	√	x
Deleting a Prometheus instance	√	x
Creating an application discovery rule	√	x
Modifying an application discovery rule	√	x
Deleting an application discovery rule	√	x
Subscribing to threshold alarms	√	x
Configuring a VM log collection path	√	x

Table 13-4 lists the common operations supported by each system-defined policy of collection settings. Select policies as required.

Table 13-4 Common operations supported by each system-defined policy of collection settings

Operation	AOM FullAccess	AOM ReadOnlyAccess
Querying a proxy area	√	√
Editing a proxy area	√	×
Deleting a proxy area	√	×
Creating a proxy area	√	×
Querying all proxies in a proxy area	√	√
Querying all proxy areas	√	√
Querying the Agent installation result	√	√

Operation	AOM FullAccess	AOM ReadOnlyAccess
Obtaining the Agent installation command of a host	√	√
Obtaining the host heartbeat and checking whether the host is connected with the server	√	√
Uninstalling running Agents in batches	√	×
Querying the Agent home page	√	√
Testing the connectivity between the installation host and the target host	√	×
Installing Agents in batches	√	×
Obtaining the latest operation log of the Agent	√	√
Obtaining the list of versions that can be selected during Agent installation	√	√
Obtaining the list of all Agent versions under the current project ID	√	√
Deleting hosts with Agents installed	√	×
Querying Agent information based on the ECS ID	√	√

Operation	AOM FullAccess	AOM ReadOnlyAccess
Deleting a host with an Agent installed	√	×
Setting an installation host	√	×
Resetting installation host parameters	√	×
Querying the list of hosts that can be set to installation hosts	√	√
Querying the list of Agent installation hosts	√	√
Deleting an installation host	√	×
Upgrading Agents in batches	√	×
Querying historical task logs	√	√
Querying historical task details	√	√
Querying all historical tasks	√	√
Querying all execution statuses and task types	√	√
Querying the Agent execution statuses in historical task details	√	√
Modifying a proxy	√	×
Deleting a proxy	√	×
Setting a proxy	√	×

Operation	AOM FullAccess	AOM ReadOnlyAccess
Querying the list of hosts that can be set to proxies	√	√
Updating plug-ins in batches	√	×
Uninstalling plug-ins in batches	√	×
Installing plug-ins in batches	√	×
Querying historical task logs of a plug-in	√	√
Querying all plug-in execution records	√	√
Querying plug-in execution records based on the task ID	√	√
Querying the plug-in execution statuses in historical task details	√	√
Obtaining the plug-in list	√	√
Querying the plug-in version	√	√
Querying the list of supported plug-ins	√	√
Obtaining the CCE cluster list	√	√
Obtaining the Agent list of a CCE cluster	√	√
Installing ICAgent on a CCE cluster	√	×
Upgrading ICAgent for a CCE cluster	√	×

Operation	AOM FullAccess	AOM ReadOnlyAccess
Uninstalling ICAgent from a CCE cluster	√	×
Obtaining the CCE cluster list	√	√
Obtaining the list of hosts where the ICAgent has been installed	√	√
Installing ICAgent on CCE cluster hosts	√	×
Upgrading ICAgent on CCE cluster hosts	√	×
Uninstalling ICAgent from CCE cluster hosts	√	×

Roles/Policies Required by AOM Dependency Services

If an IAM user needs to view data or use functions on the AOM console, grant the **AOM FullAccess** or **AOM ReadOnlyAccess** policy to the user group to which the user belongs and then add the roles or policies required by dependency services by referring to **Table 13-5**. **When a user subscribes to AOM for the first time, AOM creates an agency for the user. In addition to AOM FullAccess, the user also needs to be granted the Security Administrator permission to create and delete agencies.**

Table 13-5 Roles/Policies required by AOM dependency services

Console Function	Dependency Service	Policy/Role Required
• Workload monitoring • Cluster monitoring • Prometheus for CCE	CCE	To use workload and cluster monitoring and Prometheus for CCE, you need to set the **CCE FullAccess** and **CCE Namespace** permissions.
Data subscription	Distributed Message Service (DMS) for Kafka	To use data subscription, you need to set the **DMS ReadOnlyAccess** permission.

Console Function	Dependency Service	Policy/Role Required
• Log management • Log transfer • Log jobs • Log ingestion rules • Host group management • Log alarm rules	Log Tank Service (LTS)	To use log management, log transfer, log jobs, log ingestion rules, host group management, and log alarm rules, you need to set the **LTS FullAccess** permission. For details about the fine-grained policy permissions, see **Permissions**.
Enterprise project	Enterprise Project Management Service (EPS)	To use enterprise projects, you need to set the **EPS ReadOnlyAccess** permission. For details about the fine-grained policy permissions, see **Permissions**.

Identity Policy-based Authorization

AOM supports authorization with identity policies. **Table 13-6** lists all the system-defined identity policies for AOM. System-defined policies in identity policy-based authorization are not interoperable with those in role/policy-based authorization.

Table 13-6 System-defined identity policies for AOM

Policy Name	Description	Policy Type	Dependency Permissions	Remarks
AOMReadOnlyPolicy	Read-only permissions for AOM.	System-defined identity policy	-	After IAM5 functions are enabled, authorization for AOM will be effected. Some pages may fail to be authenticated. To solve the problem, configure **AOMFullAccessPolicy** or **AOMReadOnlyPolicy**.
AOMFullAccessPolicy	Full permissions for AOM.	System-defined identity policy	iam:agencies:createAgency and iam:agencies:deleteAgency	
AOMServiceLinkedAgencyPolicy	Agency permissions required for performing AOM cross-account operations.	System-defined identity policy	-	

Table 13-7 lists the common operations supported by each system-defined identity policy of resource monitoring. Select policies as required.

Table 13-7 Common operations supported by each system-defined identity policy of resource monitoring

Operation	AOMFullAccessPolicy	AOMReadOnlyPolicy	AOMServiceLinkedAgencyPolicy
Creating an alarm rule	√	x	√
Modifying an alarm rule	√	x	√

Operation	AOMFullAccessPolicy	AOMReadOnlyPolicy	AOMServiceLinkedAgencyPolicy
Deleting an alarm rule	√	x	√
Creating an alarm template	√	x	x
Modifying an alarm template	√	x	x
Deleting an alarm template	√	x	x
Creating an alarm action rule	√	x	√
Modifying an alarm action rule	√	x	√
Deleting an alarm action rule	√	x	√
Creating a message template	√	x	x
Modifying a message template	√	x	x
Deleting a message template	√	x	x
Creating a grouping rule	√	x	x
Modifying a grouping rule	√	x	x
Deleting a grouping rule	√	x	x
Creating a suppression rule	√	x	x

Operation	AOMFullAccessPolicy	AOMReadOnlyPolicy	AOMServiceLinkedAgencyPolicy
Modifying a suppression rule	√	x	x
Deleting a suppression rule	√	x	x
Creating a silence rule	√	x	√
Modifying a silence rule	√	x	√
Deleting a silence rule	√	x	√
Creating a dashboard	√	x	x
Modifying a dashboard	√	x	x
Deleting a dashboard	√	x	x
Creating a Prometheus instance	√	x	√
Modifying a Prometheus instance	√	x	√
Deleting a Prometheus instance	√	x	√
Creating an application discovery rule	√	x	√
Modifying an application discovery rule	√	x	√
Deleting an application discovery rule	√	x	√
Subscribing to threshold alarms	√	x	x

Operation	AOMFullAccessPolicy	AOMReadOnlyPolicy	AOMServiceLinkedAgencyPolicy
Configuring a VM log collection path	√	x	x

Table 13-8 lists the common operations supported by each system-defined identity policy of collection management.

Table 13-8 Common operations supported by each system-defined identity policy of collection management

Operation	AOMFullAccessPolicy	AOMReadOnlyPolicy
Querying a proxy area	√	√
Editing a proxy area	√	×
Deleting a proxy area	√	×
Creating a proxy area	√	×
Querying all proxies in a proxy area	√	√
Querying all proxy areas	√	√
Querying the Agent installation result	√	√
Obtaining the Agent installation command of a host	√	√
Obtaining the host heartbeat and checking whether the host is connected with the server	√	√
Uninstalling running Agents in batches	√	×
Querying the Agent home page	√	√
Testing the connectivity between the installation host and the target host	√	×

Operation	AOMFullAccessPolicy	AOMReadOnlyPolicy
Installing Agents in batches	√	×
Obtaining the latest operation log of the Agent	√	√
Obtaining the list of versions that can be selected during Agent installation	√	√
Obtaining the list of all Agent versions under the current project ID	√	√
Deleting hosts with Agents installed	√	×
Querying Agent information based on the ECS ID	√	√
Deleting a host with an Agent installed	√	×
Setting an installation host	√	×
Resetting installation host parameters	√	×
Querying the list of hosts that can be set to installation hosts	√	√
Querying the list of Agent installation hosts	√	√
Deleting an installation host	√	×
Upgrading Agents in batches	√	×
Querying historical task logs	√	√
Querying historical task details	√	√
Querying all historical tasks	√	√

Operation	AOMFullAccessPolicy	AOMReadOnlyPolicy
Querying all execution statuses and task types	√	√
Querying the Agent execution statuses in historical task details	√	√
Modifying a proxy	√	×
Deleting a proxy	√	×
Setting a proxy	√	×
Querying the list of hosts that can be set to proxies	√	√
Updating plug-ins in batches	√	×
Uninstalling plug-ins in batches	√	×
Installing plug-ins in batches	√	×
Querying historical task logs of a plug-in	√	√
Querying all plug-in execution records	√	√
Querying plug-in execution records based on the task ID	√	√
Querying the plug-in execution statuses in historical task details	√	√
Obtaining the plug-in list	√	√
Querying the plug-in version	√	√
Querying the list of supported plug-ins	√	√
Obtaining the CCE cluster list	√	√
Obtaining the Agent list of a CCE cluster	√	√

Operation	AOMFullAccessPolicy	AOMReadOnlyPolicy
Installing ICAgent on a CCE cluster	√	×
Upgrading ICAgent for a CCE cluster	√	×
Uninstalling ICAgent from a CCE cluster	√	×
Obtaining the CCE cluster list	√	√
Obtaining the list of hosts where the ICAgent has been installed	√	√
Installing ICAgent on CCE cluster hosts	√	×
Upgrading ICAgent on CCE cluster hosts	√	×
Uninstalling ICAgent from CCE cluster hosts	√	×

Identity Policy Dependencies of the AOM Console

If an IAM user needs to view data or use functions on the AOM console, grant the **AOMFullAccessPolicy** or **AOMReadOnlyAccessPolicy** permission to the user group to which the user belongs and then add the identity policies required by dependency services by referring to **Table 13-9**.

NOTE

When a user subscribes to AOM for the first time, AOM creates an agency for the user. In addition to **AOMFullAccessPolicy** or **AOMReadOnlyAccessPolicy**, the user also needs to be granted the **iam:agencies:createAgency** and **iam:agencies:deleteAgency** permissions to create and delete agencies.

Table 13-9 Identity policies required by AOM dependency services

Console Function	Dependency Service	Identity Policy Required
• Workload monitoring • Cluster monitoring	CCE	To use workload and cluster monitoring, you need to set the **CCEReadOnlyAccessPolicy** permission.

Console Function	Dependency Service	Identity Policy Required
Data subscription	Distributed Message Service (DMS) for Kafka	To use data subscription, you need to set the **DMSReadOnlyAccess** permission.
Data multi-write	• Virtual Private Cloud (VPC) • VPC Endpoint (VPCEP)	To use this function, ensure that the **VPCEndpoint ReadOnlyAccess** and **VPC ReadOnlyAccess** permissions have been set.

Reference

- **IAM Service Overview**
- **Granting Permissions Using IAM**
- **Actions Supported by Identity Policy-based Authorization**

14 Privacy Statement

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

Collector Deployment

During the installation of UniAgent and ICAgent on Linux hosts, the history recording function is disabled. Therefore, your AK/SK and access code cannot be viewed by running commands. Additionally, credentials are encrypted for storage to prevent leaks.

Container Monitoring

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

