

# CA Unified Infrastructure Management

## Probe Guide for AWS Monitoring

AWS v3.0 series



## Copyright Notice

This online help system (the "System") is for your informational purposes only and is subject to change or withdrawal by CA at any time.

This System may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This System is confidential and proprietary information of CA and protected by the copyright laws of the United States and international treaties. This System may not be disclosed by you or used for any purpose other than as may be permitted in a separate agreement between you and CA governing your use of the CA software to which the System relates (the "CA Software"). Such agreement is not modified in any way by the terms of this notice.

Notwithstanding the foregoing, if you are a licensed user of the CA Software you may make one copy of the System for internal use by you and your employees, provided that all CA copyright notices and legends are affixed to the reproduced copy.

The right to make a copy of the System is limited to the period during which the license for the CA Software remains in full force and effect. Should the license terminate for any reason, it shall be your responsibility to certify in writing to CA that all copies and partial copies of the System have been destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS SYSTEM "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS SYSTEM, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The manufacturer of this System is CA.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

Copyright © 2014 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Legal information on third-party and public domain software used in this product is documented in the *Third-Party Licenses and Terms of Use* ([http://docs.nimsoft.com/prodhelp/en\\_US/Library/Legal.html](http://docs.nimsoft.com/prodhelp/en_US/Library/Legal.html)).

# Contact CA

## Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At <http://ca.com/support>, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

## Providing Feedback about Product Documentation

Send comments or questions about CA Technologies product documentation to [nimsoft.techpubs@ca.com](mailto:nimsoft.techpubs@ca.com).

To provide feedback about general CA Technologies product documentation, complete our short customer survey which is available on the support website at <http://ca.com/docs>.



# Contents

---

**Chapter 1: Documentation Changes 7**

**Chapter 2: Overview 9**

About This Guide .....10

Related Documentation .....11

- Preconfiguration Requirements .....11
- Upgrades and Migrations .....11
- Software Requirements .....12
- NAS Subsystem ID Requirements .....12
- Alarm Threshold Requirements .....14
- Supported Platforms .....14

**Chapter 3: Configuration Details 15**

aws node .....15

- <Profile Name> node .....17
- AWS Service Health .....28

Configure a node .....29

How to Configure Alarm Thresholds .....30

Configure Static Alarm Thresholds .....30

Manage profiles .....30

Delete a profile .....31

**Chapter 4: AWS QoS Metrics 33**

**Chapter 5: Known issues 39**



# Chapter 1: Documentation Changes

---

This table describes the version history for this document.

Version	Date	What's New?
3.0	September 2014	Added description about EBS node, RDS node, ElastiCache node, and Custom Metrics node.
2.0	June 2014	Documentation on the first Admin Console GUI enabled version of the AWS Monitoring probe.



# Chapter 2: Overview

---

The Amazon Web Services (AWS) Monitoring probe remotely monitors the health and performance counters of various AWS services over a cloud network. The probe lets you create profiles that monitor your AWS user account and fetches all the service data from the AWS CloudWatch. The probe lets you configure various monitoring parameters on each service. Based on the configured parameters, the probe generates Quality of Service (QoS) data and issues status alarms.

**Note:** The probe from version 2.01 and later is configured only through the Admin Console GUI.

## **Amazon Web Service (AWS):**

The AWS provides a decentralized IT infrastructure to multiple organizations. You can create an account on the AWS cloud and can use its services as per your IT infrastructure requirements. The various capabilities of AWS include storage, web-scale computing, database access, and messaging.

The AWS Monitoring probe provides monitoring of the following AWS services:

- **Health:** The probe monitors the overall health status of the AWS services for all geographical locations. Alarms are generated based on the status of all the AWS services.
- **Amazon Simple Storage Service (S3):** This AWS service provides an interface for storing and fetching data at any time instance. The AWS Monitoring probe generates QoS data based on the time consumed in storing and retrieving files.
- **Amazon Elastic Compute Cloud (EC2):** This AWS service provides a flexible web-scale computing interface. The AWS Monitoring probe generates QoS data and alarms that are based on the performance of various EC2 instances.
- **Amazon Elastic Block Storage (EBS):** This AWS service provides a scalable storage volume facility for the EC2 instances. The AWS Monitoring probe generates QoS data and alarms that are based on the operations that are performed on the storage volumes.

- Amazon Relational Database Service (RDS): This AWS service manages relational databases that are stored in a cloud network. AWS-RDS handles many database administration tasks and lets you perform other operations like setting up and scaling the database. The AWS Monitoring probe generates QoS data and alarms that are based on the system metrics and database operations.
- Amazon ElastiCache: This AWS service provides the AWS instances with an option of storing temporary data in a scalable cache memory, and thus, increasing the processing speed. The AWS Monitoring probe generates QoS data based on the time consumed in accessing the cache service and other parameters like amount of data stored and time taken to fetch the data.
- AWS Custom Metrics: AWS provides some default metrics for all its services. Another feature of AWS is that you can create and configure your own customized metrics, and store these metrics in the AWS CloudWatch for viewing, or monitoring purpose. These metrics, which AWS does not generate, are called custom metrics. The AWS Monitoring probe lets you configure the custom metrics for QoS generation.

**Important!** Amazon charges the AWS account which the probe uses to monitor the AWS services. You must consider this fact while configuring the probe for monitoring various AWS services.

This section contains the following topics:

[About This Guide](#) (see page 10)

[Related Documentation](#) (see page 11)

## About This Guide

This guide is for the CA UIM Administrator to help understand the configuration of the AWS Monitoring probe and provides the following information:

- Overview of the AWS Monitoring probe and related documentation.
- Configuration details of the probe.
- Field description and common procedures for configuring the probe.

**Important!** Description for the intuitive GUI fields is not included in the document.

## Related Documentation

### Related Documentation

Documentation for other versions of the AWS probe

The [Release Notes](#) for the AWS probe

[User documentation for the Admin Console](#)

*Monitor Metrics Reference Information for CA Unified Infrastructure Management Probes*

([http://docs.nimsoft.com/prodhelp/en\\_US/Probes/ProbeReference/index.htm](http://docs.nimsoft.com/prodhelp/en_US/Probes/ProbeReference/index.htm))

## Preconfiguration Requirements

This section contains the preconfiguration requirements for the CA UIM AWS Monitoring probe.

- An AWS user-account with valid user-credentials, such as, **Access Key** and **Secret Access Key**.
- EC2 Administrative Rights so that the AWS Monitoring probe can access the AWS resource.

## Upgrades and Migrations

This section provides information about the upgrade and migration scenarios for the AWS Monitoring probe.

- When you install the probe version 2.01 then manually move the existing configurations, in case you are using the probe of version earlier than version 2.01.
- Delete all the versions of the AWS Monitoring probe that are earlier than version 2.01 as upgrade from a previous version to version 2.01 is not supported.
- The probe from version 2.01 and later is accessible only through the Admin Console GUI.
- For viewing the new metrics that are introduced in the AWS probe version 3.0, on the USM portal, you can perform any one of the following actions:
  - Upgrade NMS 7.6 (or earlier) to CA UIM 8.0
  - Install the **ci\_defn\_pack** version 1.00 probe. you are required to restart the **nis\_server** when you deploy the **ci\_defn\_pack**.

**Important!** You can install the **ci\_defn\_pack** probe from <https://support.nimsoft.com>

## Software Requirements

The probe requires:

- One, or more Amazon AWS user-accounts, and EC2 administrative privilege.

The AWS probe version 2.01 and later has the following requirements along with the above mentioned requirements:

- NMS version 7.6, or CA UIM version 8.0 and later.

## NAS Subsystem ID Requirements

Alarms are classified by their subsystem ID, identifying which part of the system the alarm relates to. These subsystem IDs are kept in a table maintained by the NAS probe. If you are working with NMS 7.6 or earlier, you will have to add the following subsystem IDs manually using the NAS Raw Configuration menu. However, if you have upgraded to CA UIM 8.0 then you do not have to manually add the following subsystem IDs:

Key Name	Value
2.19.	Amazon
2.19.1.	AWS
2.19.1.1.	Resource
2.19.1.2.	ServiceStatus
2.19.1.3.	EC2
2.19.1.4.	S3
2.19.1.5.	EBS
2.19.1.6.	RDS
2.19.1.9	ElastiCache

**To update the Subsystem IDs using Admin Console, follow these steps:**

1. In the Admin Console, click the black arrow next to the NAS probe, select **Raw Configure**.
2. Click on the **Subsystems** folder.
3. Click on the New Key Menu item.
4. Enter the Key Name in the Add key window, click **Add**.  
The new key appears in the list of keys with a blank value.
5. Click in the Value column for the newly created key and enter the key value.
6. Repeat this process for all of the required subsystem IDs for your probe.
7. Click **Apply**.

**To update the Subsystem IDs using Infrastructure Manager, follow these steps:**

1. In Infrastructure Manager, right click on the NAS probe, select **Raw Configure**.
2. Click on the **Subsystems** folder.
3. Click on the **New Key...** button.
4. Enter the Key Name and Value, Click **OK**.
5. Repeat this process for all of the required subsystem IDs for your probe.
6. Click **Apply**.

## Alarm Threshold Requirements

The PPM probe maintains a table of subsystem IDs that are mapped to the probes. As of the current release, the subsystem IDs for this probe will default to 1.1.19. The AWS probe supports the following types of alarms:

- Dynamic Alarms Thresholds
- Static Alarm Thresholds
- Time-over Thresholds
- Time-to Thresholds

If you are using either dynamic or static alarm thresholds, you can change the default entry to the appropriate subsystem ID.

**Note:** If you have upgraded NMS 7.6 to CA UIM 8.0 then you do not have to follow the following procedure.

**Follow these steps:**

1. In the Admin Console, click the black arrow next to the probe, select **Configure**.
2. Select the monitor that you want to modify from the available list.
3. Change the Static and Dynamic **Subsystem (override) fields** to **2.19.1.1**.
4. Save your settings.

## Supported Platforms

Refer to the CA UIM [Compatibility Support Matrix](#) for the latest information on supported platforms. See also the [Support Matrix for CA UIM Probes](#) for additional specific information on the AWS probe.

# Chapter 3: Configuration Details

---

The AWS Monitoring probe is configured to create monitoring profiles for accessing AWS resources and fetching data from AWS CloudWatch. You can also configure health monitors to generate alarms on the basis of the availability of services in various geographical regions.

The probe also lets you configure the **Auto Discovery** functionality. If any service instance is added or deleted in the AWS resource, then the **Auto Discovery** functionality updates the list of instances in the probe.

The probe fetches data of instances, or services and provides you with various monitors for generating QoS. You can also configure the probe to fetch the list of custom metrics that are created for a specific service in the AWS CloudWatch.

This section contains the following topics:

[aws node](#) (see page 15)

[Configure a node](#) (see page 29)

[How to Configure Alarm Thresholds](#) (see page 30)

[Configure Static Alarm Thresholds](#) (see page 30)

[Manage profiles](#) (see page 30)

[Delete a profile](#) (see page 31)

## aws node

This node lets you view the probe information and configure the logging properties. You can also set the polling interval for **Auto Discovery** functionality and configure the proxy settings.

**Note:** The AWS services nodes are visible in the **Navigation Pane** only after you create a monitoring profile. Initially, only the **AWS** node and the **AWS Service Health** node are visible.

**Navigation:** aws

Set or modify the following values as required:

**aws > Probe Information**

This section provides information about the probe name, probe version, start time of the probe, and the probe vendor.

**aws > Probe Setup**

This section lets you configure the detail level of the log file. The default value is 3-info.

**aws > Auto Discovery**

This section lets you set the value of **Discovery Interval (minutes)**. If any instance is added or deleted in the AWS resource, then the **Auto Discovery** functionality updates the list of instances in the probe. The **Discovery Interval (minutes)** specifies the time interval between each time the probe runs the Auto Discovery functionality.

**aws > Proxy Settings**

This section enables you to connect to the AWS cloud through a proxy server on the network. You need proxy server settings when your network is not an open network.

- Enable Proxy: lets you use a proxy server for connecting to the AWS cloud.
- IP: defines the IP address of the proxy server.
- Port: specifies the port on the proxy server through which the connection is established.
- Username: defines the user name for accessing the proxy server.

### aws > Add New Profile

This section lets you add a profile for monitoring the AWS services data. QoS data is generated according to the performance of these services.

- **Profile Name:** defines a unique name for the monitoring profile. This field was identified as **Name** in the previous versions of the probe. You can specify the AWS account name as the value for this field.
- **Active:** activates the profile for service monitoring.
- **Interval (seconds):** specifies the time interval (in seconds) after which the probe collects the data from the AWS cloud for the specific profile.

Default: 600

**Note:** Interval value must be above 300 seconds.

- **Alarm Message:** specifies the alarm to be generated when the connection to AWS services fails.

Default: ResourceCritical

- **Access Key:** defines the login credential of the AWS user-account for accessing the AWS resource.
- **Secret Access Key:** specifies the additional login credential of the AWS user-account.

**Note:** The probe uses the combination of the **Access Key** and **Secret Access Key** for accessing the AWS resource.

## <Profile Name> node

This node represents the profile which is created to monitor the health and performance of AWS services. Each profile is mapped with an AWS account. You can check the connection between the probe and the AWS resource through the **Verify Credentials** button under the **Actions** drop down.

**Note:** This node is referred to as *profile name* node in the document and is user-configurable.

**Navigation:** AWS > *profile name*

Refer to the **Add New Profile** section of the **AWS node** topic for field description.

## EC2 node

The AWS EC2 service of a specific region stores the instance data in AWS CloudWatch. For a specific profile, the AWS Monitoring probe fetches the data from AWS CloudWatch.

This node lets you configure the probe for interacting with the EC2 service and collect data about the instances of the AWS resource. The probe generates QoS based on the instance data which is collected from AWS CloudWatch.

**Navigation:** AWS > *profile name* > EC2

Set or modify the following values as required:

### EC2 > EC2 Configurations

This node lets you configure EC2 service properties.

- **Active:** activates the addition of custom metrics from the AWS CloudWatch.
- **Start Time:** specifies the time duration (in minutes) for collecting sample values from the AWS CloudWatch. The probe starts collecting the values that were calculated during the time period which is specified here.
- **Statistics:** defines one of the following operations to be performed on the sample values that the probe fetches:
  - Calculate minimum value.
  - Calculate maximum value.
  - Calculate the sum of all the values.
  - Calculate the average of all the values.

Default: Average

**Note:** When you change the Statistics value, the QoS graphs on the UMP portal get changed.

- **Period (minutes):** specifies a time interval which is used to divide the collected values into groups.

For example, if the **Start Time** is specified as 10 minutes and the **Period** is specified as 2 minutes, then the values are fetched for 5 minutes time interval.

## <Instance Name> node

This node represents an instance of the AWS resource. An EC2 instance is a virtual machine (VM). If any region subscribes to the EC2 service, then an instance of EC2 VM is created for that region.

The AWS Monitoring probe monitors the performance counters of the EC2 instances of the AWS resource. All EC2 instances are visible under the EC2 node.

**Note:** This node is referred to as *instance name* node in the document and each instance has a unique ID.

This node does not contain any fields or sections.

## <Monitor Name> node

This node lets you configure the performance counters of the EC2 instances. The AWS Monitoring probe generates QoS data of the EC2 service of a specific region according to the values fetched from the AWS CloudWatch.

The performance counters are divided into following categories:

- CPU
- Disk
- Network

Each category is represented as a node under the *instance name* node.

**Note:** This node is referred to as *EC2-monitor name* node in the document and it represents various EC2 performance counters.

**Navigation:** AWS > *profile name* > EC2 > *instance name* > *EC2-monitor name*

Set or modify the following values as required:

### ***monitor name* > Monitors**

This section lets you configure the performance counters for generating QoS.

**Note:** The performance counters of an EC2 instance are visible in a tabular form. You can select any one counter in the table and can configure its properties.

- QoS Name: indicates the name of performance counter.
- Publish Data: generates the QoS data for the selected counter.

**Note:** When you select the **Publish Data** check box, the value of the **Data** column in the table changes from **Off** to **On**.

Similarly, you can configure the other performance counters that are visible under the **CPU**, **Disk**, and **Network** nodes.

## <EBS Volume> node

This node represents the Elastic Block Storage (EBS) which is linked to a specific EC2 instance. The EBS node is visible in the navigation panel only if you have added a storage block with the EC2 instances, or when you have assigned a storage block to the instances.

**Note:** This node is referred to as *EBS Volume* node in the document and it represents an EBS storage volume.

**Navigation:** AWS > *profile name* > EC2 > *instance name* > EBS

Set or modify the following values as required:

### **EBS > Monitors**

This section lets you configure the performance counters for generating QoS.

**Note:** The performance counters of an EBS volume are visible in a tabular form. You can select any one counter in the table and can configure its properties.

- QoS Name: indicates the name of performance counter.
- Publish Data: generates the QoS data for the selected counter.

**Note:** When you select the **Publish Data** check box, the value of the **Data** column in the table changes from **Off** to **On**.

## ElastiCache node

The AWS ElastiCache service provides a scalable cache for storing temporary data. The AWS Monitoring probe generates QoS based on the instance data which is collected from CloudWatch.

This node lets you configure the probe to fetch ElastiCache instance information.

**Navigation:** AWS > *profile name* > ElastiCache

Set or modify the following values as required:

### **ElastiCache > ElastiCache Configurations**

This node lets you configure the ElastiCache service properties. For field descriptions, refer to the **EC2** topic.

## <Instance Name> node

This node represents an AWS instance that uses the ElastiCache service. The ElastiCache service supports two types of cache engines:

- Remote Dictionary Server or Redis (Currently, ElastiCache supports a single-node Redis cache cluster)
- Memcached (Currently, ElastiCache supports a maximum of 20 nodes in a cache cluster)

The instances are displayed in the navigation pane according to the type of cache engine.

**Note:** This node is known as *instance name* node in the document and each instance has a unique ID.

This node does not contain any fields or sections.

## <Node Name> node

This node lets you configure the performance counters of the ElastiCache instances. The AWS probe generates QoS data of the ElastiCache service according to the values fetched from CloudWatch.

**Note:** This node is referred to as *node name* node in the document and it represents a node of a Memcached or Redis ElastiCache instance.

**Navigation:** AWS > *profile name* > Elasti Cache > *instance name* > *node name*

Set or modify the following values as required:

### **monitor name** > Monitors

This section lets you configure the performance counters for generating QoS.

**Note:** The performance counters of an ElastiCache instance are visible in a tabular form. You can select any one counter in the table and can configure its properties.

- QoS Name: indicates the name of performance counter.
- Publish Data: generates the QoS data for the selected counter.

**Note:** When you select the **Publish Data** check box, the value of the **Data** column in the table changes from **Off** to **On**.

<ElastiCache-Monitor Name> node

This node lets you configure the performance counters of an ElastiCache instance node.

The performance counters are divided into following categories:

- CPU
- Memory

Each category is represented as a node under the *node name* node.

**Note:** This node is referred to as *ElastiCache-monitor name* node in the document and it represents various ElastiCache instance performance counters.

**Navigation:** AWS > *profile name* > ElastiCache > *instance name* > *node name* > *ElastiCache-monitor name*

Set or modify the following values as required:

#### **ElastiCache-monitor name > Monitors**

This section lets you configure the performance counters for generating QoS.

**Note:** The performance counters of an ElastiCache instance node are visible in a tabular form. You can select any one counter in the table and can configure its properties.

- QoS Name: indicates the name of performance counter.
- Publish Data: generates the QoS data for the selected counter.

**Note:** When you select the **Publish Data** check box, the value of the **Data** column in the table changes from **Off** to **On**.

Similarly, you can configure the other performance counters that are visible under the **Memory** node.

## RDS node

The AWS RDS service manages relational databases that are stored on AWS CloudWatch. The AWS Monitoring probe fetches the data from the CloudWatch and generates QoS related to an RDS instance.

**Navigation:** AWS > *profile name* > RDS

Set or modify the following values as required:

#### **RDS > RDS Configurations**

This node lets you configure the RDS service properties.

- Active: activates the addition of database instances of the AWS resource. For field descriptions, refer to the **EC2** topic.

## <Database Name> node

**Note:** This node is known as *database name* node in the document.

**Navigation:** AWS > *profile name* > RDS > *database name*

Set or modify the following values as required:

### *database name* > Monitors

This section lets you configure the performance counters of a relational database instance for generating QoS data.

**Note:** The performance counters of an RDS database instance are visible in a tabular form. You can select any one counter in the table and can configure its properties.

- QoS Name: indicates the name of performance counter.
- Publish Data: generates the QoS data for the selected counter.

**Note:** When you select the **Publish Data** check box, the value of the **Data** column in the table changes from **Off** to **On**.

## <RDS Monitor Name> node

This node lets you configure the performance counters of RDS instances. The AWS probe generates QoS data of the RDS service according to the values fetched from AWS CloudWatch.

The performance counters are divided into following categories:

- CPU
- Disk
- Memory
- Network

Each category is represented as a node under the *database name* node.

**Note:** This node is referred to as *RDS monitor name* node in the document and it represents various RDS performance counters.

**Navigation:** AWS > *profile name* > RDS > *database name* > *RDS monitor name*

Set or modify the following values as required:

### **RDS monitor name > Monitors**

This section lets you configure the RDS performance counters of a specific instance for generating QoS data.

**Note:** The performance counters of a relational database are visible in a tabular form. You can select any one counter in the table and can configure its properties.

- QoS Name: indicates the name of performance counter.
- Publish Data: generates the QoS data for the selected counter.

**Note:** When you select the **Publish Data** check box, the value of the **Data** column in the table changes from **Off** to **On**.

Similarly, you can configure the other performance counters that are visible under the **CPU**, **Disk**, **Memory**, and **Network** nodes.

## Custom Metrics

In AWS, metrics are segregated into different **Namespaces**. A **Dimension** is a variable that categorizes a metric according to its statistics. When you create custom metrics through the script and store the metrics in AWS CloudWatch, the AWS probe fetches that data from CloudWatch.

This node lets you select a custom metric that is available in an AWS Namespace and then define custom QoS for it. The custom metrics for different AWS Namespace are visible in the **Navigation Pane**.

You can configure any of the discovered metrics that are available in an AWS Namespace through the **Custom Metrics** node except RDS, EC2, EBS, and ElastiCache.

**Navigation:** AWS > *profile name* > Custom Metric

Set or modify the following values as required:

### Custom Metric > Custom Configurations

This section lets you configure the probe to fetch the list of custom metrics from the AWS CloudWatch and select custom metrics for a specific Namespace.

- **Available Service Metrics:** specifies the list of available AWS Namespaces that the probe fetches from CloudWatch. Each Namespace contains various custom metrics. You can move specific service Namespace from the **Available List** to the **Selected List**. The selected service metrics are visible as nodes in the **Navigation Pane**.

**Note:** For other field descriptions, refer to the **EC2** topic.

## <AWS-Service Name> node

This node lets you view and configure the custom metrics for all AWS services. You can define a custom QoS name, unit, and can let the probe generate QoS data for the custom metric. This node contains a table that lists the AWS dimensions against each service metric.

**Note:** This node is referred to as *AWS-service name* node in the document and is user-configurable.

**Navigation:** AWS > *profile name* > Custom Metric > *AWS-service name*

Set or modify the following values as required:

### **AWS-service name > Collected Metrics**

This section lets you define custom QoS name for different service metrics that are listed in a tabular form. You can also configure the probe to generate QoS data for selected metrics.

**Note:** If you have created the custom metrics in a custom Namespace then only custom metrics are visible in the table. However, if you have created the custom metrics in an existing Namespace then all the metrics are visible in the table.

- Publish Data: generates the QoS data for the selected counter.

**Note:** When you select the **Publish Data** check box, the value of the **Data** column in the table changes from **Off** to **On**.

- QoS Name: defines a custom name for the selected service metric. The identifier **QoS** is automatically prefixed with the specified QoS name. So, for example, if you give the custom QoS name as **abc**, then the metric name is generated as **QoS\_abc**.
- QoS Unit: defines the unit for measuring the value of the defined QoS.  
**Note:** Ensure that the QoS unit is the same as it is in the script which generates the QoS data.
- Dimension: indicates a unique identifier for the selected metric.

## S3 node

The data which is stored in the cloud using the AWS S3 service is segregated into groups that are known as buckets. The AWS probe monitors the time which is consumed in storing and retrieving files to and from the bucket, respectively.

This node lets you configure the performance counters for S3 service. The AWS probe generates QoS data related to the time that is consumed in storing and retrieving files to and from the S3 buckets.

**Note:** Set the polling interval (**Interval** field in the **Add New Profile** section in **aws** node) according to the size of the file that you want to store or retrieve. If the polling interval is too less, then the probe starts fetching data again from the bucket before completing a previous file process. For example, if you want to upload a file of size 1 MB then you can set the polling interval as 5 minutes.

**Navigation:** *AWS > profile name > S3*

Set or modify the following values as required:

### S3 > S3 Configurations

This section lets you provide details about the file bucket so that the probe can monitor the time that is consumed in accessing the file bucket.

- **Active:** enables the monitoring of file bucket access time.
- **Bucket Name:** specifies the name of the file bucket for which the probe monitors the storing and retrieving time.
- **File Name:** defines the name of the file which is stored or retrieved from the bucket.

**Note:** The file, for which you want to generate the QoS data, must be present in the AWS probe base folder (**/probes/applications/aws**).

### S3 > Monitors

This section lets you configure the performance counters for generating QoS.

**Note:** The performance counters of the S3 service are visible in a tabular form. You can select any one counter in the table and can configure its properties.

- **QoS Name:** indicates the name of performance counter.
- **Publish Data:** generates the QoS data for the selected counter.

**Note:** When you select the **Publish Data** check box, the value of the **Data** column in the table changes from **Off** to **On**.

## AWS Service Health

This node represents the health monitoring service of AWS probe. The probe monitors AWS service availability for a specific region. The probe generates alarms in case any service for a specific region is unavailable. The following alarms are generated after the probe monitors the health of the AWS services:

- Disruption in the service.
- Performance issues.
- Service is operating normally.
- Other information.

**Navigation:** AWS > AWS Service Health

Set or modify the following values as required:

### AWS Service Health > Health Configuration

This section enables you to configure the Health Monitoring functionality of the AWS probe. The **Health Interval (mins)** field lets you set the time interval, in minutes, during which the probe fetches the health status of the AWS services.

## <AWS Region> node

This node lets you view the list of AWS services that are available for a specific region. You can configure the AWS probe for generating alarms for specific AWS services in a region.

**Note:** This node is known as *AWS region* in the document as this node represents all the geographical locations where AWS provides services.

**Navigation:** AWS > AWS Service Health > *AWS region*

Set or modify the following values as required:

**AWS Region > AWS Service Status**

This section lets you view the various AWS services that are available for a specific region. You can configure the service properties for generating the alarms in case the service is not available.

**Note:** The AWS services for the selected region are visible in a tabular form. You can select any one service in the table and can configure its properties.

- Description: indicates the description of the selected service.
- Unit: indicates the unit of the selected service status.
- Metric Type ID: identifies a unique ID for alarm generation.
- Publish Alarms: enables the probe to check the status of the selected service and generate alarms.

**Note:** When you select the **Publish Alarms** check box, the value of the **Alarm** column in the table changes from **Off** to **On**.

- Service: indicates the name of the selected service.

Similarly, you can configure the services of the other geographical locations.

## Configure a node

This procedure provides the information to configure a section within a node.

Each section within the node lets you configure the properties of the probe for connecting to the AWS resource and monitoring various AWS services.

**Follow these steps:**

1. Navigate to the section within a node that you want to configure.
2. Update the field information and click **Save**.

The specified section of the probe is configured.

## How to Configure Alarm Thresholds

Some Quality of Service measurement probes allow you to set different types of alarm thresholds. These threshold options allow you to more broadly control when alarm messages are sent for each QoS probe.

For more information about the different alarm thresholds and their configuration requirements, refer to the *General Probe Configuration* section of the Admin Console Help.

## Configure Static Alarm Thresholds

**Important!** In order to create static alarm thresholds, you must have the `baseline_engine probe version 2.2` installed on the robot and configured.

Static thresholds can be set at the QoS metric level in some of the probes that publish alarms for a QoS metric. For more information about static alarm thresholds, refer to the *How to Configure Static Thresholds* section of the Admin Console Help.

## Manage profiles

The following procedure enables you to add a profile for monitoring the AWS services. Each profile represents one AWS resource. There can be multiple instances of an AWS resource.

**Follow these steps:**

1. Click **Options** next to the **AWS** node in the navigation pane.
2. Select **Add New Profile**.
3. Update the field information and click **Submit**.

The new monitoring profile is visible under the AWS node in the navigation pane.

The **Auto Discovery** functionality automatically loads a list of all the available instances.

## Delete a profile

You can delete a profile if you do not want the probe to monitor the performance of a specific AWS resource.

**Follow these steps:**

1. Click the **Options** icon next to the *profile name* node that you want to delete.
2. Select **Delete Profile**.
3. Click **Save**.

The monitoring profile is deleted from the resource.



# Chapter 4: AWS QoS Metrics

---

The following table describes the checkpoint metrics that can be configured using the AWS Monitoring probe.

The following QoS data is for the AWS S3 service:

QoS Name	Metric Name	Units	Description
QOS_AWS_FILEREADTIME	File Read Time	Seconds	Time taken to fetch a file from the file bucket.
QOS_AWS_FILEWRITETIME	File Write Time	Seconds	Time taken to store a file in the file bucket.

The following QoS data is for the AWS EC2 service:

QoS Name	Metric Name	Units	Description
QOS_AWS_CPU_UTILIZATION	CPU Usage	Percent	The percentage of allocated EC2 compute units that are currently in use on the instance. This metric identifies the processing power required to run an application upon a selected instance.
QOS_AWS_DISK_WRITE_BYTES	Data Written	Bytes	This metric is used to determine the volume of the data the application writes onto the hard disk of the instance. This can be used to determine the speed of the application.
QOS_AWS_DISK_READ_BYTES	Data Read	Bytes	This metric is used to determine the volume of the data the application reads from the hard disk of the instance. This can be used to determine the speed of the application.
QOS_AWS_DISK_READ_OPS	Reads	Count	Completed read operations from all ephemeral disks available to the instance. This metric identifies the rate at which an application reads a disk. This can be used to determine the speed in which an application reads data from a hard disk.
QOS_AWS_DISK_WRITE_OPS	Writes	Count	Completed write operations to all ephemeral disks available to the instance. This metric identifies the rate at which an application writes to a hard disk. This can be used to determine the speed in which an application saves data to a hard disk.

QOS_AWS_NETWORK_IN	Total Bytes Received	Bytes	The number of bytes received on all network interfaces by the instance. This metric identifies the volume of incoming network traffic to an application on a single instance.
QOS_AWS_NETWORK_OUT	Total Bytes Sent	Bytes	The number of bytes sent out on all network interfaces by the instance. This metric identifies the volume of outgoing network traffic to an application on a single instance.

The following QoS data is for the AWS EBS service:

QoS Name	Metric Name	Units	Description
QOS_AWS_VOLUME_READ_BYTES	Total Read	Bytes	The number of bytes read in the time period specified in the <b>Start Time</b> field.
QOS_AWS_VOLUME_WRITE_BYTES	Total Written	Bytes	The number of bytes written in the time period specified in the <b>Start Time</b> field..
QOS_AWS_VOLUME_READ_OPS	Total Read Operations	Count	The number of Read operations in the time period specified in the <b>Start Time</b> field.
QOS_AWS_VOLUME_WRITE_OPS	Total Write Operations	Count	The number of Write operations in the time period specified in the <b>Start Time</b> field.
QOS_AWS_VOLUME_TOTAL_READ_TIME	Total Read Time	Seconds	The number of seconds spent by all read operations that completed in a specified period of time. If multiple requests are submitted at the same time, this total could be greater than the length of the period.
QOS_AWS_VOLUME_TOTAL_WRITE_TIME	Total Write Time	Seconds	The number of seconds spent by all write operations that completed in a specified period of time. If multiple requests are submitted at the same time, this total could be greater than the length of the period.
QOS_AWS_VOLUME_IDLE_TIME	Total Idle Time	Seconds	The number of seconds in the time period specified in the <b>Start Time</b> field when no read or write operations were submitted.
QOS_AWS_VOLUME_QUEUE_LENGTH	Queue Length	Count	The number of read and write operation requests waiting to be completed in the time period specified in the <b>Start Time</b> field.

QOS_AWS_VOLUME_THROUGHPUT_PERCENTAGE	Throughput Percentage	Percentage	The percentage of I/O operations per second (IOPS) delivered of the total IOPS provisioned for an Amazon EBS volume. Used with Provisioned IOPS (SSD) volumes only.
QOS_AWS_VOLUME_CONSUMED_READ_WRITE_OPS	Consumed Read Write Operations	Count	The total amount of read and write operations (normalized to 16K capacity units) consumed in the time period specified in the <b>Start Time</b> field. Used with Provisioned IOPS (SSD) volumes only.

The following QoS data is for the AWS RDS service:

QoS Name	Metric Name	Units	Description
QOS_AWS_RDS_BIN_LOG_DISK_USAGE	Binary Log Size On Disk	Bytes	The amount of disk space occupied by binary logs on the master. Applies to MySQL read replicas only.
QOS_AWS_RDS_CPU_UTILIZATION	CPU Utilization	Percent	The percentage of CPU utilization by the relational database.
QOS_AWS_RDS_DATABASE_CONNECTIONS	Database Connections	Count	The number of database connections in use.
QOS_AWS_RDS_DISK_QUEUE_DEPTH	Outstanding IOs in queue	Count	The number of outstanding IOs (read/write requests) waiting to access the disk.
QOS_AWS_RDS_FREEABLE_MEMORY	Available Memory	Bytes	The amount of available Random Access Memory.
QOS_AWS_RDS_FREE_STORAGE_SPACE	Available Storage Space	Bytes	The amount of available storage space.
QOS_AWS_RDS_REPLICA_LAG	Read Replica Lag Time	Seconds	The time for which a <b>Read Replica DB</b> instance lags behind the <b>Source DB</b> instance. Replica lag occurs due to slow execution of data manipulation queries and is seen in MySQL database only.
QOS_AWS_RDS_SWAP_USAGE	Used Swap Space	Bytes	The amount of swap space used on the DB Instance.

QOS_AWS_RDS_READ_IOPS	Read Operations Per Second	Count/Second	The average number of disk I/O operations per second.
QOS_AWS_RDS_WRITE_IOPS	Write Operations Per Second	Count/Second	The average number of disk I/O operations per second.
QOS_AWS_RDS_READ_LATENCY	Read Latency	Seconds	The average amount of time taken per disk I/O operation.
QOS_AWS_RDS_WRITE_LATENCY	Write Latency	Seconds	The average amount of time taken per disk I/O operation.
QOS_AWS_RDS_READ_THROUGHPUT	Read Throughput	Bytes/Second	The average number of bytes read from disk per second.
QOS_AWS_RDS_WRITE_THROUGHPUT	Write Throughput	Bytes/Second	The average number of bytes written to disk per second.
QOS_AWS_RDS_NETWORK_RECEIVE_THROUGHPUT	Network Receive Throughput	Bytes	The incoming (Receive) network traffic on the DB instance, including both customer database traffic and Amazon RDS traffic used for monitoring and replication.
QOS_AWS_RDS_NETWORK_TRANSMIT_THROUGHPUT	Network Transmit Throughput	Bytes	The outgoing (Transmit) network traffic on the DB instance, including both customer database traffic and Amazon RDS traffic used for monitoring and replication.

The following QoS data is for the AWS ElastiCache service:

#### Host Level Metrics:

QoS Name	Metric Name	Units	Description
QOS_AWS_ELASTICACHE_CPU_UTILIZATION	CPU Utilization	Percent	The percentage of CPU utilization.
QOS_AWS_ELASTICACHE_FREEABLE_MEMORY	Available Free Memory	Bytes	The amount of free memory available on the host.

**Memcached Metrics:**

QoS Name	Metric Name	Units	Description
QOS_AWS_ELASTICACHE_MEMCACHED_UNUSED_MEMORY	Unused Memory For Cache	Bytes	The amount of unused memory the cache can use to store items. This is derived from the memcached statistics limit_maxbytes and bytes by subtracting bytes from limit_maxbytes.
QOS_AWS_ELASTICACHE_MEMCACHED_CURRENT_ITEMS	Number of Items	Count	A count of the number of items currently stored in the cache.
QOS_AWS_ELASTICACHE_MEMCACHED_EVICTIONS	Total Non-Expired Evicted Items	Count	The number of non-expired items the cache evicted to allow space for new writes.
QOS_AWS_ELASTICACHE_MEMCACHED_RECLAIMED	Total Expired Items Evicted	Count	The number of expired items the cache evicted to allow space for new writes
QOS_AWS_ELASTICACHE_MEMCACHED_GET_HITS	Total Cache Hits Requests	Count	The number of get requests the cache has received where the key requested was found.
QOS_AWS_ELASTICACHE_MEMCACHED_GET_MISSES	Total Cache Miss Requests	Count	The number of get requests the cache has received where the key requested was not found
QOS_AWS_ELASTICACHE_MEMCACHED_BYTES_USED_FOR_CACHE_ITEMS	Total Bytes Used for Cache Items	Bytes	The number of bytes used to store cache items.

**Redis Metrics:**

QoS Name	Metric Name	Units	Description
QOS_AWS_ELASTICACHE_REDIS_CURRENT_CONNECTIONS	Total Bytes Allocated for Cache	Count	The number of client connections, excluding connections from read replicas.

QOS_AWS_ELASTICCACHE_REDIS_BYTES_USED_FOR_CACHE	Number of Connections	Bytes	The number of bytes allocated by Redis.
---	-----------------------	-------	---

# Chapter 5: Known issues

---

This section contains a list of known issues in this release:

- Each time that you save the configuration, the AWS Monitoring probe is restarted and data collection for AWS services, starts again. This re-discovery of services slows the GUI processing. You are required to re-load the GUI after the probe restarts.