



Welcome

Application Signals



API Version 2024-04-15

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Application Signals: Welcome

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Welcome

Actions

BatchGetServiceLevelObjectiveBudgetReport

Request Syntax

URI Request Parameters

Request Body

Response Syntax

Response Elements

Errors

See Also

BatchUpdateExclusionWindows

Request Syntax

URI Request Parameters

Request Body

Response Syntax

Response Elements

Errors

See Also

CreateServiceLevelObjective

Request Syntax

URI Request Parameters

Request Body

Response Syntax

Response Elements

Errors

See Also

DeleteServiceLevelObjective

Request Syntax

URI Request Parameters

Request Body

Response Syntax

Response Elements

Errors

See Also

1

2

3

3

3

3

4

7

7

8

9

9

9

10

10

11

11

12

13

14

17

17

19

21

22

22

24

24

24

24

24

24

24

25

GetService	26
Request Syntax	26
URI Request Parameters	26
Request Body	26
Response Syntax	27
Response Elements	28
Errors	29
See Also	30
GetServiceLevelObjective	31
Request Syntax	31
URI Request Parameters	31
Request Body	31
Response Syntax	31
Response Elements	34
Errors	34
See Also	35
ListServiceDependencies	36
Request Syntax	36
URI Request Parameters	36
Request Body	37
Response Syntax	38
Response Elements	38
Errors	39
See Also	40
ListServiceDependents	41
Request Syntax	41
URI Request Parameters	41
Request Body	42
Response Syntax	43
Response Elements	43
Errors	44
See Also	45
ListServiceLevelObjectiveExclusionWindows	46
Request Syntax	46
URI Request Parameters	46
Request Body	46

Response Syntax	46
Response Elements	47
Errors	47
See Also	48
ListServiceLevelObjectives	49
Request Syntax	49
URI Request Parameters	49
Request Body	50
Response Syntax	51
Response Elements	52
Errors	52
See Also	53
ListServiceOperations	54
Request Syntax	54
URI Request Parameters	54
Request Body	55
Response Syntax	56
Response Elements	56
Errors	57
See Also	58
ListServices	59
Request Syntax	59
URI Request Parameters	59
Request Body	60
Response Syntax	60
Response Elements	61
Errors	62
See Also	62
ListTagsForResource	63
Request Syntax	63
URI Request Parameters	63
Request Body	63
Response Syntax	63
Response Elements	64
Errors	64
See Also	64

StartDiscovery	66
Request Syntax	66
URI Request Parameters	66
Request Body	66
Response Syntax	66
Response Elements	66
Errors	67
See Also	67
TagResource	68
Request Syntax	68
URI Request Parameters	68
Request Body	68
Response Syntax	69
Response Elements	69
Errors	69
See Also	70
UntagResource	71
Request Syntax	71
URI Request Parameters	71
Request Body	71
Response Syntax	72
Response Elements	72
Errors	72
See Also	72
UpdateServiceLevelObjective	74
Request Syntax	74
URI Request Parameters	76
Request Body	77
Response Syntax	78
Response Elements	80
Errors	81
See Also	81
Data Types	82
BatchUpdateExclusionWindowsError	84
Contents	84
See Also	84

BurnRateConfiguration	86
Contents	86
See Also	86
CalendarInterval	87
Contents	87
See Also	88
DependencyConfig	89
Contents	89
See Also	90
Dimension	91
Contents	91
See Also	91
ExclusionWindow	93
Contents	93
See Also	94
Goal	95
Contents	95
See Also	96
Interval	97
Contents	97
See Also	97
Metric	98
Contents	98
See Also	98
MetricDataQuery	100
Contents	100
See Also	102
MetricReference	104
Contents	104
See Also	105
MetricStat	106
Contents	106
See Also	107
MonitoredRequestCountMetricDataQueries	108
Contents	108
See Also	108

RecurrenceRule	110
Contents	110
See Also	110
RequestBasedServiceLevelIndicator	111
Contents	111
See Also	111
RequestBasedServiceLevelIndicatorConfig	113
Contents	113
See Also	113
RequestBasedServiceLevelIndicatorMetric	115
Contents	115
See Also	117
RequestBasedServiceLevelIndicatorMetricConfig	118
Contents	118
See Also	120
RollingInterval	121
Contents	121
See Also	121
Service	122
Contents	122
See Also	124
ServiceDependency	125
Contents	125
See Also	126
ServiceDependent	127
Contents	127
See Also	128
ServiceLevelIndicator	129
Contents	129
See Also	129
ServiceLevelIndicatorConfig	131
Contents	131
See Also	131
ServiceLevelIndicatorMetric	133
Contents	133
See Also	134

ServiceLevelIndicatorMetricConfig	135
Contents	135
See Also	137
ServiceLevelObjective	138
Contents	138
See Also	140
ServiceLevelObjectiveBudgetReport	142
Contents	142
See Also	145
ServiceLevelObjectiveBudgetReportError	146
Contents	146
See Also	147
ServiceLevelObjectiveSummary	148
Contents	148
See Also	150
ServiceOperation	151
Contents	151
See Also	151
ServiceSummary	152
Contents	152
See Also	154
Tag	155
Contents	155
See Also	155
Window	156
Contents	156
See Also	156
Common Parameters	157
Common Errors	160

Welcome

Use CloudWatch Application Signals for comprehensive observability of your cloud-based applications. It enables real-time service health dashboards and helps you track long-term performance trends against your business goals. The application-centric view provides you with unified visibility across your applications, services, and dependencies, so you can proactively monitor and efficiently triage any issues that may arise, ensuring optimal customer experience.

Application Signals provides the following benefits:

- Automatically collect metrics and traces from your applications, and display key metrics such as call volume, availability, latency, faults, and errors.
- Create and monitor service level objectives (SLOs).
- See a map of your application topology that Application Signals automatically discovers, that gives you a visual representation of your applications, dependencies, and their connectivity.

Application Signals works with CloudWatch RUM, CloudWatch Synthetics canaries, and AWS Service Catalog AppRegistry, to display your client pages, Synthetics canaries, and application names within dashboards and maps.

This document was last published on September 3, 2025.

Actions

The following actions are supported:

- [BatchGetServiceLevelObjectiveBudgetReport](#)
- [BatchUpdateExclusionWindows](#)
- [CreateServiceLevelObjective](#)
- [DeleteServiceLevelObjective](#)
- [GetService](#)
- [GetServiceLevelObjective](#)
- [ListServiceDependencies](#)
- [ListServiceDependents](#)
- [ListServiceLevelObjectiveExclusionWindows](#)
- [ListServiceLevelObjectives](#)
- [ListServiceOperations](#)
- [ListServices](#)
- [ListTagsForResource](#)
- [StartDiscovery](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateServiceLevelObjective](#)

BatchGetServiceLevelObjectiveBudgetReport

Use this operation to retrieve one or more *service level objective (SLO) budget reports*.

An *error budget* is the amount of time or requests in an unhealthy state that your service can accumulate during an interval before your overall SLO budget health is breached and the SLO is considered to be unmet. For example, an SLO with a threshold of 99.95% and a monthly interval translates to an error budget of 21.9 minutes of downtime in a 30-day month.

Budget reports include a health indicator, the attainment value, and remaining budget.

For more information about SLO error budgets, see [SLO concepts](#).

Request Syntax

```
POST /budget-report HTTP/1.1
Content-type: application/json
```

```
{
  "SloIds": [ "string" ],
  "Timestamp": number
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

SloIds

An array containing the IDs of the service level objectives that you want to include in the report.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 50 items.

Required: Yes

Timestamp

The date and time that you want the report to be for. It is expressed as the number of milliseconds since Jan 1, 1970 00:00:00 UTC.

Type: Timestamp

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "Errors": [
    {
      "Arn": "string",
      "ErrorCode": "string",
      "ErrorMessage": "string",
      "Name": "string"
    }
  ],
  "Reports": [
    {
      "Arn": "string",
      "Attainment": number,
      "BudgetRequestsRemaining": number,
      "BudgetSecondsRemaining": number,
      "BudgetStatus": "string",
      "EvaluationType": "string",
      "Goal": {
        "AttainmentGoal": number,
        "Interval": { ... },
        "WarningThreshold": number
      },
      "Name": "string",
      "RequestBasedSli": {
        "ComparisonOperator": "string",
        "MetricThreshold": number,
        "RequestBasedSliMetric": {
          "DependencyConfig": {
            "DependencyKeyAttributes": {
```

```

        "string" : "string"
      },
      "DependencyOperationName": "string"
    },
    "KeyAttributes": {
      "string" : "string"
    },
    "MetricType": "string",
    "MonitoredRequestCountMetric": { ... },
    "OperationName": "string",
    "TotalRequestCountMetric": [
      {
        "AccountId": "string",
        "Expression": "string",
        "Id": "string",
        "Label": "string",
        "MetricStat": {
          "Metric": {
            "Dimensions": [
              {
                "Name": "string",
                "Value": "string"
              }
            ],
            "MetricName": "string",
            "Namespace": "string"
          },
          "Period": number,
          "Stat": "string",
          "Unit": "string"
        },
        "Period": number,
        "ReturnData": boolean
      }
    ]
  },
  "Sli": {
    "ComparisonOperator": "string",
    "MetricThreshold": number,
    "SliMetric": {
      "DependencyConfig": {
        "DependencyKeyAttributes": {
          "string" : "string"
        }
      }
    }
  }
}

```

```

    },
    "DependencyOperationName": "string"
  },
  "KeyAttributes": {
    "string" : "string"
  },
  "MetricDataQueries": [
    {
      "AccountId": "string",
      "Expression": "string",
      "Id": "string",
      "Label": "string",
      "MetricStat": {
        "Metric": {
          "Dimensions": [
            {
              "Name": "string",
              "Value": "string"
            }
          ],
          "MetricName": "string",
          "Namespace": "string"
        },
        "Period": number,
        "Stat": "string",
        "Unit": "string"
      },
      "Period": number,
      "ReturnData": boolean
    }
  ],
  "MetricType": "string",
  "OperationName": "string"
}
},
"TotalBudgetRequests": number,
"TotalBudgetSeconds": number
}
],
"Timestamp": number
}

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Errors

An array of structures, where each structure includes an error indicating that one of the requests in the array was not valid.

Type: Array of [ServiceLevelObjectiveBudgetReportError](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Reports

An array of structures, where each structure is one budget report.

Type: Array of [ServiceLevelObjectiveBudgetReport](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Timestamp

The date and time that the report is for. It is expressed as the number of milliseconds since Jan 1, 1970 00:00:00 UTC.

Type: Timestamp

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

BatchUpdateExclusionWindows

Add or remove time window exclusions for one or more Service Level Objectives (SLOs).

Request Syntax

```
PATCH /exclusion-windows HTTP/1.1
Content-type: application/json

{
  "AddExclusionWindows": [
    {
      "Reason": "string",
      "RecurrenceRule": {
        "Expression": "string"
      },
      "StartTime": number,
      "Window": {
        "Duration": number,
        "DurationUnit": "string"
      }
    }
  ],
  "RemoveExclusionWindows": [
    {
      "Reason": "string",
      "RecurrenceRule": {
        "Expression": "string"
      },
      "StartTime": number,
      "Window": {
        "Duration": number,
        "DurationUnit": "string"
      }
    }
  ],
  "SloIds": [ "string" ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

AddExclusionWindows

A list of exclusion windows to add to the specified SLOs. You can add up to 10 exclusion windows per SLO.

Type: Array of [ExclusionWindow](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Required: No

RemoveExclusionWindows

A list of exclusion windows to remove from the specified SLOs. The window configuration must match an existing exclusion window.

Type: Array of [ExclusionWindow](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Required: No

SloIds

The list of SLO IDs to add or remove exclusion windows from.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 50 items.

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "Errors": [
```

```
{
  {
    "ErrorCode": "string",
    "ErrorMessage": "string",
    "SloId": "string"
  },
  "SloIds": [ "string" ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Errors

A list of errors that occurred while processing the request.

Type: Array of [BatchUpdateExclusionWindowsError](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

SloIds

The list of SLO IDs that were successfully processed.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 50 items.

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ResourceNotFoundException

Resource not found.

HTTP Status Code: 404

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

CreateServiceLevelObjective

Creates a service level objective (SLO), which can help you ensure that your critical business operations are meeting customer expectations. Use SLOs to set and track specific target levels for the reliability and availability of your applications and services. SLOs use service level indicators (SLIs) to calculate whether the application is performing at the level that you want.

Create an SLO to set a target for a service or operation's availability or latency. CloudWatch measures this target frequently you can find whether it has been breached.

The target performance quality that is defined for an SLO is the *attainment goal*.

You can set SLO targets for your applications that are discovered by Application Signals, using critical metrics such as latency and availability. You can also set SLOs against any CloudWatch metric or math expression that produces a time series.

Note

You can't create an SLO for a service operation that was discovered by Application Signals until after that operation has reported standard metrics to Application Signals.

When you create an SLO, you specify whether it is a *period-based SLO* or a *request-based SLO*. Each type of SLO has a different way of evaluating your application's performance against its attainment goal.

- A *period-based SLO* uses defined *periods* of time within a specified total time interval. For each period of time, Application Signals determines whether the application met its goal. The attainment rate is calculated as the number of good periods/number of total periods.

For example, for a period-based SLO, meeting an attainment goal of 99.9% means that within your interval, your application must meet its performance goal during at least 99.9% of the time periods.

- A *request-based SLO* doesn't use pre-defined periods of time. Instead, the SLO measures number of good requests/number of total requests during the interval. At any time, you can find the ratio of good requests to total requests for the interval up to the time stamp that you specify, and measure that ratio against the goal set in your SLO.

After you have created an SLO, you can retrieve error budget reports for it. An *error budget* is the amount of time or amount of requests that your application can be non-compliant with the SLO's goal, and still have your application meet the goal.

- For a period-based SLO, the error budget starts at a number defined by the highest number of periods that can fail to meet the threshold, while still meeting the overall goal. The *remaining error budget* decreases with every failed period that is recorded. The error budget within one interval can never increase.

For example, an SLO with a threshold that 99.95% of requests must be completed under 2000ms every month translates to an error budget of 21.9 minutes of downtime per month.

- For a request-based SLO, the remaining error budget is dynamic and can increase or decrease, depending on the ratio of good requests to total requests.

For more information about SLOs, see [Service level objectives \(SLOs\)](#).

When you perform a `CreateServiceLevelObjective` operation, Application Signals creates the `AWSServiceRoleForCloudWatchApplicationSignals` service-linked role, if it doesn't already exist in your account. This service-linked role has the following permissions:

- `xray:GetServiceGraph`
- `logs:StartQuery`
- `logs:GetQueryResults`
- `cloudwatch:GetMetricData`
- `cloudwatch:ListMetrics`
- `tag:GetResources`
- `autoscaling:DescribeAutoScalingGroups`

Request Syntax

```
POST /slo HTTP/1.1
Content-type: application/json

{
  "BurnRateConfigurations": [
    {
      "LookBackWindowMinutes": number
```

```
    }
  ],
  "Description": "string",
  "Goal": {
    "AttainmentGoal": number,
    "Interval": { ... },
    "WarningThreshold": number
  },
  "Name": "string",
  "RequestBasedSliConfig": {
    "ComparisonOperator": "string",
    "MetricThreshold": number,
    "RequestBasedSliMetricConfig": {
      "DependencyConfig": {
        "DependencyKeyAttributes": {
          "string" : "string"
        },
        "DependencyOperationName": "string"
      },
      "KeyAttributes": {
        "string" : "string"
      },
      "MetricType": "string",
      "MonitoredRequestCountMetric": { ... },
      "OperationName": "string",
      "TotalRequestCountMetric": [
        {
          "AccountId": "string",
          "Expression": "string",
          "Id": "string",
          "Label": "string",
          "MetricStat": {
            "Metric": {
              "Dimensions": [
                {
                  "Name": "string",
                  "Value": "string"
                }
              ],
              "MetricName": "string",
              "Namespace": "string"
            },
            "Period": number,
            "Stat": "string",
```



```
        "Unit": "string"
      },
      "Period": number,
      "ReturnData": boolean
    }
  ]
}
},
"SliConfig": {
  "ComparisonOperator": "string",
  "MetricThreshold": number,
  "SliMetricConfig": {
    "DependencyConfig": {
      "DependencyKeyAttributes": {
        "string": "string"
      },
      "DependencyOperationName": "string"
    },
    "KeyAttributes": {
      "string": "string"
    },
    "MetricDataQueries": [
      {
        "AccountId": "string",
        "Expression": "string",
        "Id": "string",
        "Label": "string",
        "MetricStat": {
          "Metric": {
            "Dimensions": [
              {
                "Name": "string",
                "Value": "string"
              }
            ],
            "MetricName": "string",
            "Namespace": "string"
          },
          "Period": number,
          "Stat": "string",
          "Unit": "string"
        },
        "Period": number,
        "ReturnData": boolean
      }
    ]
  }
}
```

```
    }
  ],
  "MetricType": "string",
  "OperationName": "string",
  "PeriodSeconds": number,
  "Statistic": "string"
}
},
"Tags": [
  {
    "Key": "string",
    "Value": "string"
  }
]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

BurnRateConfigurations

Use this array to create *burn rates* for this SLO. Each burn rate is a metric that indicates how fast the service is consuming the error budget, relative to the attainment goal of the SLO.

Type: Array of [BurnRateConfiguration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Required: No

Description

An optional description for this SLO.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Goal

This structure contains the attributes that determine the goal of the SLO.

Type: [Goal](#) object

Required: No

Name

A name for this SLO.

Type: String

Pattern: `[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]`

Required: Yes

RequestBasedSliConfig

If this SLO is a request-based SLO, this structure defines the information about what performance metric this SLO will monitor.

You can't specify both `RequestBasedSliConfig` and `SliConfig` in the same operation.

Type: [RequestBasedServiceLevelIndicatorConfig](#) object

Required: No

SliConfig

If this SLO is a period-based SLO, this structure defines the information about what performance metric this SLO will monitor.

You can't specify both `RequestBasedSliConfig` and `SliConfig` in the same operation.

Type: [ServiceLevelIndicatorConfig](#) object

Required: No

Tags

A list of key-value pairs to associate with the SLO. You can associate as many as 50 tags with an SLO. To be able to associate tags with the SLO when you create the SLO, you must have the `cloudwatch:TagResource` permission.

Tags can help you organize and categorize your resources. You can also use them to scope user permissions by granting a user permission to access or change only resources with certain tag values.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "Slo": {
    "Arn": "string",
    "BurnRateConfigurations": [
      {
        "LookBackWindowMinutes": number
      }
    ],
    "CreatedTime": number,
    "Description": "string",
    "EvaluationType": "string",
    "Goal": {
      "AttainmentGoal": number,
      "Interval": { ... },
      "WarningThreshold": number
    },
    "LastUpdatedTime": number,
    "MetricSourceType": "string",
    "Name": "string",
    "RequestBasedSli": {
      "ComparisonOperator": "string",
      "MetricThreshold": number,
      "RequestBasedSliMetric": {
        "DependencyConfig": {
          "DependencyKeyAttributes": {
            "string" : "string"
          }
        },
```

```

    "DependencyOperationName": "string"
  },
  "KeyAttributes": {
    "string" : "string"
  },
  "MetricType": "string",
  "MonitoredRequestCountMetric": { ... },
  "OperationName": "string",
  "TotalRequestCountMetric": [
    {
      "AccountId": "string",
      "Expression": "string",
      "Id": "string",
      "Label": "string",
      "MetricStat": {
        "Metric": {
          "Dimensions": [
            {
              "Name": "string",
              "Value": "string"
            }
          ],
          "MetricName": "string",
          "Namespace": "string"
        },
        "Period": number,
        "Stat": "string",
        "Unit": "string"
      },
      "Period": number,
      "ReturnData": boolean
    }
  ]
},
"SLI": {
  "ComparisonOperator": "string",
  "MetricThreshold": number,
  "SLIMetric": {
    "DependencyConfig": {
      "DependencyKeyAttributes": {
        "string" : "string"
      },
      "DependencyOperationName": "string"
    }
  }
}

```

```

    },
    "KeyAttributes": {
      "string" : "string"
    },
    "MetricDataQueries": [
      {
        "AccountId": "string",
        "Expression": "string",
        "Id": "string",
        "Label": "string",
        "MetricStat": {
          "Metric": {
            "Dimensions": [
              {
                "Name": "string",
                "Value": "string"
              }
            ],
            "MetricName": "string",
            "Namespace": "string"
          },
          "Period": number,
          "Stat": "string",
          "Unit": "string"
        },
        "Period": number,
        "ReturnData": boolean
      }
    ],
    "MetricType": "string",
    "OperationName": "string"
  }
}

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Slo

A structure that contains information about the SLO that you just created.

Type: [ServiceLevelObjective](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You don't have sufficient permissions to perform this action.

HTTP Status Code: 403

ConflictException

This operation attempted to create a resource that already exists.

HTTP Status Code: 409

ServiceQuotaExceededException

This request exceeds a service quota.

HTTP Status Code: 402

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteServiceLevelObjective

Deletes the specified service level objective.

Request Syntax

```
DELETE /slo/Id HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

Id

The ARN or name of the service level objective to delete.

Pattern: `[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]$|`

`^arn:aws:application-signals:[^:]*:[^:]*:slo/[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ResourceNotFoundException

Resource not found.

HTTP Status Code: 404

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetService

Returns information about a service discovered by Application Signals.

Request Syntax

```
POST /service?EndTime=EndTime&StartTime=StartTime HTTP/1.1
Content-type: application/json
```

```
{
  "KeyAttributes": {
    "string" : "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

EndTime

The end of the time period to retrieve information about. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

Your requested start time will be rounded to the nearest hour.

Required: Yes

StartTime

The start of the time period to retrieve information about. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

Your requested start time will be rounded to the nearest hour.

Required: Yes

Request Body

The request accepts the following data in JSON format.

KeyAttributes

Use this field to specify which service you want to retrieve information for. You must specify at least the Type, Name, and Environment attributes.

This is a string-to-string map. It can include the following fields.

- `Type` designates the type of object this is.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.
- `Identifier` identifies the resource objects of this resource. This is used only if the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Environment` specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json
```

```
{
  "EndTime": number,
  "LogGroupReferences": [
    {
      "string" : "string"
    }
  ]
}
```

```
],
  "Service": {
    "AttributeMaps": [
      {
        "string" : "string"
      }
    ],
    "KeyAttributes": {
      "string" : "string"
    },
    "LogGroupReferences": [
      {
        "string" : "string"
      }
    ],
    "MetricReferences": [
      {
        "AccountId": "string",
        "Dimensions": [
          {
            "Name": "string",
            "Value": "string"
          }
        ],
        "MetricName": "string",
        "MetricType": "string",
        "Namespace": "string"
      }
    ]
  },
  "StartTime": number
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

EndTime

The end time of the data included in the response. In a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057.

This displays the time that Application Signals used for the request. It might not match your request exactly, because it was rounded to the nearest hour.

Type: Timestamp

[LogGroupReferences](#)

An array of string-to-string maps that each contain information about one log group associated with this service. Each string-to-string map includes the following fields:

- "Type": "AWS::Resource"
- "ResourceType": "AWS::Logs::LogGroup"
- "Identifier": "*name-of-log-group*"

Type: Array of string to string maps

Map Entries: Maximum number of 4 items.

Key Pattern: [a-zA-Z]{1,50}

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: [-~]*[!-~]+[-~]*

[Service](#)

A structure containing information about the service.

Type: [Service](#) object

[StartTime](#)

The start time of the data included in the response. In a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057.

This displays the time that Application Signals used for the request. It might not match your request exactly, because it was rounded to the nearest hour.

Type: Timestamp

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetServiceLevelObjective

Returns information about one SLO created in the account.

Request Syntax

```
GET /slo/Id HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

Id

The ARN or name of the SLO that you want to retrieve information about. You can find the ARNs of SLOs by using the [ListServiceLevelObjectives](#) operation.

Pattern: `[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]$|`

`^arn:aws:application-signals:[^:]*:[^:]*:slo/[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "Slo": {
    "Arn": "string",
    "BurnRateConfigurations": [
      {
        "LookBackWindowMinutes": number
      }
    ],
    "CreatedTime": number,
```



```
"Description": "string",
"EvaluationType": "string",
"Goal": {
  "AttainmentGoal": number,
  "Interval": { ... },
  "WarningThreshold": number
},
"LastUpdatedTime": number,
"MetricSourceType": "string",
"Name": "string",
"RequestBasedSli": {
  "ComparisonOperator": "string",
  "MetricThreshold": number,
  "RequestBasedSliMetric": {
    "DependencyConfig": {
      "DependencyKeyAttributes": {
        "string" : "string"
      },
      "DependencyOperationName": "string"
    },
    "KeyAttributes": {
      "string" : "string"
    },
    "MetricType": "string",
    "MonitoredRequestCountMetric": { ... },
    "OperationName": "string",
    "TotalRequestCountMetric": [
      {
        "AccountId": "string",
        "Expression": "string",
        "Id": "string",
        "Label": "string",
        "MetricStat": {
          "Metric": {
            "Dimensions": [
              {
                "Name": "string",
                "Value": "string"
              }
            ],
            "MetricName": "string",
            "Namespace": "string"
          },
          "Period": number,
```

```

        "Stat": "string",
        "Unit": "string"
    },
    "Period": number,
    "ReturnData": boolean
}
]
}
},
"SLI": {
    "ComparisonOperator": "string",
    "MetricThreshold": number,
    "SLIMetric": {
        "DependencyConfig": {
            "DependencyKeyAttributes": {
                "string" : "string"
            },
            "DependencyOperationName": "string"
        },
        "KeyAttributes": {
            "string" : "string"
        },
        "MetricDataQueries": [
            {
                "AccountId": "string",
                "Expression": "string",
                "Id": "string",
                "Label": "string",
                "MetricStat": {
                    "Metric": {
                        "Dimensions": [
                            {
                                "Name": "string",
                                "Value": "string"
                            }
                        ],
                        "MetricName": "string",
                        "Namespace": "string"
                    },
                    "Period": number,
                    "Stat": "string",
                    "Unit": "string"
                },
                "Period": number,

```

```
        "ReturnData": boolean
      }
    ],
    "MetricType": "string",
    "OperationName": "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Slo

A structure containing the information about the SLO.

Type: [ServiceLevelObjective](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ResourceNotFoundException

Resource not found.

HTTP Status Code: 404

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListServiceDependencies

Returns a list of service dependencies of the service that you specify. A dependency is an infrastructure component that an operation of this service connects with. Dependencies can include AWS services, AWS resources, and third-party services.

Request Syntax

```
POST /service-dependencies?
EndTime=EndTime&MaxResults=MaxResults&NextToken=NextToken&StartTime=StartTime HTTP/1.1
Content-type: application/json

{
  "KeyAttributes": {
    "string" : "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

EndTime

The end of the time period to retrieve information about. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

Your requested end time will be rounded to the nearest hour.

Required: Yes

MaxResults

The maximum number of results to return in one operation. If you omit this parameter, the default of 50 is used.

Valid Range: Minimum value of 1. Maximum value of 100.

NextToken

Include this value, if it was returned by the previous operation, to get the next set of service dependencies.

StartTime

The start of the time period to retrieve information about. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

Your requested start time will be rounded to the nearest hour.

Required: Yes

Request Body

The request accepts the following data in JSON format.

KeyAttributes

Use this field to specify which service you want to retrieve information for. You must specify at least the Type, Name, and Environment attributes.

This is a string-to-string map. It can include the following fields.

- `Type` designates the type of object this is.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.
- `Identifier` identifies the resource objects of this resource. This is used only if the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Environment` specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "EndTime": number,
  "NextToken": "string",
  "ServiceDependencies": [
    {
      "DependencyKeyAttributes": {
        "string" : "string"
      },
      "DependencyOperationName": "string",
      "MetricReferences": [
        {
          "AccountId": "string",
          "Dimensions": [
            {
              "Name": "string",
              "Value": "string"
            }
          ],
          "MetricName": "string",
          "MetricType": "string",
          "Namespace": "string"
        }
      ],
      "OperationName": "string"
    }
  ],
  "StartTime": number
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

EndTime

The end of the time period that the returned information applies to. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

This displays the time that Application Signals used for the request. It might not match your request exactly, because it was rounded to the nearest hour.

Type: Timestamp

NextToken

Include this value in your next use of this API to get next set of service dependencies.

Type: String

ServiceDependencies

An array, where each object in the array contains information about one of the dependencies of this service.

Type: Array of [ServiceDependency](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

StartTime

The start of the time period that the returned information applies to. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

This displays the time that Application Signals used for the request. It might not match your request exactly, because it was rounded to the nearest hour.

Type: Timestamp

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListServiceDependents

Returns the list of dependents that invoked the specified service during the provided time range. Dependents include other services, CloudWatch Synthetics canaries, and clients that are instrumented with CloudWatch RUM app monitors.

Request Syntax

```
POST /service-dependents?
EndTime=EndTime&MaxResults=MaxResults&NextToken=NextToken&StartTime=StartTime HTTP/1.1
Content-type: application/json

{
  "KeyAttributes": {
    "string" : "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

[EndTime](#)

The end of the time period to retrieve information about. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

Your requested start time will be rounded to the nearest hour.

Required: Yes

[MaxResults](#)

The maximum number of results to return in one operation. If you omit this parameter, the default of 50 is used.

Valid Range: Minimum value of 1. Maximum value of 100.

[NextToken](#)

Include this value, if it was returned by the previous operation, to get the next set of service dependents.

StartTime

The start of the time period to retrieve information about. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

Your requested start time will be rounded to the nearest hour.

Required: Yes

Request Body

The request accepts the following data in JSON format.

KeyAttributes

Use this field to specify which service you want to retrieve information for. You must specify at least the Type, Name, and Environment attributes.

This is a string-to-string map. It can include the following fields.

- `Type` designates the type of object this is.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.
- `Identifier` identifies the resource objects of this resource. This is used only if the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Environment` specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "EndTime": number,
  "NextToken": "string",
  "ServiceDependents": [
    {
      "DependentKeyAttributes": {
        "string" : "string"
      },
      "DependentOperationName": "string",
      "MetricReferences": [
        {
          "AccountId": "string",
          "Dimensions": [
            {
              "Name": "string",
              "Value": "string"
            }
          ],
          "MetricName": "string",
          "MetricType": "string",
          "Namespace": "string"
        }
      ],
      "OperationName": "string"
    }
  ],
  "StartTime": number
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

EndTime

The end of the time period that the returned information applies to. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

This displays the time that Application Signals used for the request. It might not match your request exactly, because it was rounded to the nearest hour.

Type: Timestamp

NextToken

Include this value in your next use of this API to get next set of service dependents.

Type: String

ServiceDependents

An array, where each object in the array contains information about one of the dependents of this service.

Type: Array of [ServiceDependent](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

StartTime

The start of the time period that the returned information applies to. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

This displays the time that Application Signals used for the request. It might not match your request exactly, because it was rounded to the nearest hour.

Type: Timestamp

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListServiceLevelObjectiveExclusionWindows

Retrieves all exclusion windows configured for a specific SLO.

Request Syntax

```
GET /slo/Id/exclusion-windows?MaxResults=MaxResults&NextToken=NextToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

Id

The ID of the SLO to list exclusion windows for.

Pattern: `[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]$|
^arn:aws:application-signals:[^:]*:[^:]*:slo/[0-9A-Za-z][-._0-9A-Za-z]
{0,126}[0-9A-Za-z]`

Required: Yes

MaxResults

The maximum number of results to return in one operation. If you omit this parameter, the default of 50 is used.

Valid Range: Minimum value of 1. Maximum value of 10.

NextToken

Include this value, if it was returned by the previous operation, to get the next set of service level objectives.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Content-type: application/json

```
{
  "ExclusionWindows": [
    {
      "Reason": "string",
      "RecurrenceRule": {
        "Expression": "string"
      },
      "StartTime": number,
      "Window": {
        "Duration": number,
        "DurationUnit": "string"
      }
    }
  ],
  "NextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

ExclusionWindows

A list of exclusion windows configured for the SLO.

Type: Array of [ExclusionWindow](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

NextToken

Include this value, if it was returned by the previous operation, to get the next set of service level objectives.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ResourceNotFoundException

Resource not found.

HTTP Status Code: 404

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListServiceLevelObjectives

Returns a list of SLOs created in this account.

Request Syntax

```
POST /slos?
IncludeLinkedAccounts=IncludeLinkedAccounts&MaxResults=MaxResults&NextToken=NextToken&OperationName=OperationName
HTTP/1.1
Content-type: application/json

{
  "DependencyConfig": {
    "DependencyKeyAttributes": {
      "string" : "string"
    },
    "DependencyOperationName": "string"
  },
  "KeyAttributes": {
    "string" : "string"
  },
  "MetricSourceTypes": [ "string" ]
}
```

URI Request Parameters

The request uses the following URI parameters.

IncludeLinkedAccounts

If you are using this operation in a monitoring account, specify `true` to include SLO from source accounts in the returned data.

When you are monitoring an account, you can use AWS account ID in `KeyAttribute` filter for service source account and `SloOwnerawsaccountID` for SLO source account with `IncludeLinkedAccounts` to filter the returned data to only a single source account.

MaxResults

The maximum number of results to return in one operation. If you omit this parameter, the default of 50 is used.

Valid Range: Minimum value of 1. Maximum value of 50.

NextToken

Include this value, if it was returned by the previous operation, to get the next set of service level objectives.

OperationName

The name of the operation that this SLO is associated with.

Length Constraints: Minimum length of 1. Maximum length of 255.

SloOwnerAwsAccountId

SLO's AWS account ID.

Pattern: [0-9]{12}

Request Body

The request accepts the following data in JSON format.

DependencyConfig

Identifies the dependency using the `DependencyKeyAttributes` and `DependencyOperationName`.

Type: [DependencyConfig](#) object

Required: No

KeyAttributes

You can use this optional field to specify which services you want to retrieve SLO information for.

This is a string-to-string map. It can include the following fields.

- `Type` designates the type of object this is.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.

- **Identifier** identifies the resource objects of this resource. This is used only if the value of the **Type** field is **Resource** or **AWS::Resource**.
- **Environment** specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: No

MetricSourceTypes

Use this optional field to only include SLOs with the specified metric source types in the output. Supported types are:

- Service operation
- Service dependency
- CloudWatch metric

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 3 items.

Valid Values: `ServiceOperation` | `CloudWatchMetric` | `ServiceDependency`

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "NextToken": "string",
  "SloSummaries": [
    {
      "Arn": "string",
```

```
"CreatedTime": number,
"DependencyConfig": {
  "DependencyKeyAttributes": {
    "string" : "string"
  },
  "DependencyOperationName": "string"
},
"EvaluationType": "string",
"KeyAttributes": {
  "string" : "string"
},
"MetricSourceType": "string",
"Name": "string",
"OperationName": "string"
}
]
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

NextToken

Include this value in your next use of this API to get next set of service level objectives.

Type: String

SloSummaries

An array of structures, where each structure contains information about one SLO.

Type: Array of [ServiceLevelObjectiveSummary](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListServiceOperations

Returns a list of the *operations* of this service that have been discovered by Application Signals. Only the operations that were invoked during the specified time range are returned.

Request Syntax

```
POST /service-operations?
EndTime=EndTime&MaxResults=MaxResults&NextToken=NextToken&StartTime=StartTime HTTP/1.1
Content-type: application/json

{
  "KeyAttributes": {
    "string" : "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

EndTime

The end of the time period to retrieve information about. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

Your requested end time will be rounded to the nearest hour.

Required: Yes

MaxResults

The maximum number of results to return in one operation. If you omit this parameter, the default of 50 is used.

Valid Range: Minimum value of 1. Maximum value of 100.

NextToken

Include this value, if it was returned by the previous operation, to get the next set of service operations.

StartTime

The start of the time period to retrieve information about. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

Your requested start time will be rounded to the nearest hour.

Required: Yes

Request Body

The request accepts the following data in JSON format.

KeyAttributes

Use this field to specify which service you want to retrieve information for. You must specify at least the Type, Name, and Environment attributes.

This is a string-to-string map. It can include the following fields.

- Type designates the type of object this is.
- ResourceType specifies the type of the resource. This field is used only when the value of the Type field is Resource or AWS::Resource.
- Name specifies the name of the object. This is used only if the value of the Type field is Service, RemoteService, or AWS::Service.
- Identifier identifies the resource objects of this resource. This is used only if the value of the Type field is Resource or AWS::Resource.
- Environment specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: [a-zA-Z]{1,50}

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: [-~]*[!-~]+[-~]*

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "EndTime": number,
  "NextToken": "string",
  "ServiceOperations": [
    {
      "MetricReferences": [
        {
          "AccountId": "string",
          "Dimensions": [
            {
              "Name": "string",
              "Value": "string"
            }
          ],
          "MetricName": "string",
          "MetricType": "string",
          "Namespace": "string"
        }
      ],
      "Name": "string"
    }
  ],
  "StartTime": number
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

EndTime

The end of the time period that the returned information applies to. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

This displays the time that Application Signals used for the request. It might not match your request exactly, because it was rounded to the nearest hour.

Type: Timestamp

NextToken

Include this value in your next use of this API to get next set of service operations.

Type: String

ServiceOperations

An array of structures that each contain information about one operation of this service.

Type: Array of [ServiceOperation](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

StartTime

The start of the time period that the returned information applies to. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

This displays the time that Application Signals used for the request. It might not match your request exactly, because it was rounded to the nearest hour.

Type: Timestamp

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListServices

Returns a list of services that have been discovered by Application Signals. A service represents a minimum logical and transactional unit that completes a business function. Services are discovered through Application Signals instrumentation.

Request Syntax

```
GET /services?  
AwsAccountId=AwsAccountId&EndTime=EndTime&IncludeLinkedAccounts=IncludeLinkedAccounts&MaxResults=MaxResults  
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

AwsAccountId

AWS Account ID.

Pattern: [0-9]{12}

EndTime

The end of the time period to retrieve information about. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

Your requested start time will be rounded to the nearest hour.

Required: Yes

IncludeLinkedAccounts

If you are using this operation in a monitoring account, specify `true` to include services from source accounts in the returned data.

MaxResults

The maximum number of results to return in one operation. If you omit this parameter, the default of 50 is used.

Valid Range: Minimum value of 1. Maximum value of 100.

NextToken

Include this value, if it was returned by the previous operation, to get the next set of services.

StartTime

The start of the time period to retrieve information about. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

Your requested start time will be rounded to the nearest hour.

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "EndTime": number,
  "NextToken": "string",
  "ServiceSummaries": [
    {
      "AttributeMaps": [
        {
          "string" : "string"
        }
      ],
      "KeyAttributes": {
        "string" : "string"
      },
      "MetricReferences": [
        {
          "AccountId": "string",
          "Dimensions": [
            {
              "Name": "string",
              "Value": "string"
            }
          ]
        }
      ]
    }
  ]
}
```

```
    ],  
    "MetricName": "string",  
    "MetricType": "string",  
    "Namespace": "string"  
  }  
]  
}  
],  
"StartTime": number  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

EndTime

The end of the time period that the returned information applies to. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

This displays the time that Application Signals used for the request. It might not match your request exactly, because it was rounded to the nearest hour.

Type: Timestamp

NextToken

Include this value in your next use of this API to get next set of services.

Type: String

ServiceSummaries

An array of structures, where each structure contains some information about a service. To get complete information about a service, use [GetService](#).

Type: Array of [ServiceSummary](#) objects

StartTime

The start of the time period that the returned information applies to. When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: 1698778057

This displays the time that Application Signals used for the request. It might not match your request exactly, because it was rounded to the nearest hour.

Type: Timestamp

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListTagsForResource

Displays the tags associated with a CloudWatch resource. Tags can be assigned to service level objectives.

Request Syntax

```
GET /tags?ResourceArn=ResourceArn HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

ResourceArn

The Amazon Resource Name (ARN) of the CloudWatch resource that you want to view tags for.

The ARN format of an Application Signals SLO is `arn:aws:cloudwatch:Region:account-id:slo:slo-name`

For more information about ARN format, see [Resource Types Defined by Amazon CloudWatch](#) in the *Amazon Web Services General Reference*.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "Tags": [
    {
      "Key": "string",
```



```
    "Value": "string"  
  }  
]  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Tags

The list of tag keys and values associated with the resource you specified.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ResourceNotFoundException

Resource not found.

HTTP Status Code: 404

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)

- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

StartDiscovery

Enables this AWS account to be able to use CloudWatch Application Signals by creating the *AWSServiceRoleForCloudWatchApplicationSignals* service-linked role. This service-linked role has the following permissions:

- xray:GetServiceGraph
- logs:StartQuery
- logs:GetQueryResults
- cloudwatch:GetMetricData
- cloudwatch:ListMetrics
- tag:GetResources
- autoscaling:DescribeAutoScalingGroups

After completing this step, you still need to instrument your Java and Python applications to send data to Application Signals. For more information, see [Enabling Application Signals](#).

Request Syntax

```
POST /start-discovery HTTP/1.1
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You don't have sufficient permissions to perform this action.

HTTP Status Code: 403

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

TagResource

Assigns one or more tags (key-value pairs) to the specified CloudWatch resource, such as a service level objective.

Tags can help you organize and categorize your resources. You can also use them to scope user permissions by granting a user permission to access or change only resources with certain tag values.

Tags don't have any semantic meaning to AWS and are interpreted strictly as strings of characters.

You can use the TagResource action with an alarm that already has tags. If you specify a new tag key for the alarm, this tag is appended to the list of tags associated with the alarm. If you specify a tag key that is already associated with the alarm, the new tag value that you specify replaces the previous value for that tag.

You can associate as many as 50 tags with a CloudWatch resource.

Request Syntax

```
POST /tag-resource HTTP/1.1
Content-type: application/json
```

```
{
  "ResourceArn": "string",
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

ResourceArn

The Amazon Resource Name (ARN) of the CloudWatch resource that you want to set tags for.

The ARN format of an Application Signals SLO is `arn:aws:cloudwatch:Region:account-id:slo:slo-name`

For more information about ARN format, see [Resource Types Defined by Amazon CloudWatch](#) in the *Amazon Web Services General Reference*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

Tags

The list of key-value pairs to associate with the alarm.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ResourceNotFoundException

Resource not found.

HTTP Status Code: 404

ServiceQuotaExceededException

This request exceeds a service quota.

HTTP Status Code: 402

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UntagResource

Removes one or more tags from the specified resource.

Request Syntax

```
POST /untag-resource HTTP/1.1
Content-type: application/json
```

```
{
  "ResourceArn": "string",
  "TagKeys": [ "string" ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[ResourceArn](#)

The Amazon Resource Name (ARN) of the CloudWatch resource that you want to delete tags from.

The ARN format of an Application Signals SLO is `arn:aws:cloudwatch:Region:account-id:slo:slo-name`

For more information about ARN format, see [Resource Types Defined by Amazon CloudWatch](#) in the *Amazon Web Services General Reference*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

[TagKeys](#)

The list of tag keys to remove from the resource.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Length Constraints: Minimum length of 1. Maximum length of 128.

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ResourceNotFoundException

Resource not found.

HTTP Status Code: 404

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UpdateServiceLevelObjective

Updates an existing service level objective (SLO). If you omit parameters, the previous values of those parameters are retained.

You cannot change from a period-based SLO to a request-based SLO, or change from a request-based SLO to a period-based SLO.

Request Syntax

```
PATCH /slo/Id HTTP/1.1
Content-type: application/json

{
  "BurnRateConfigurations": [
    {
      "LookBackWindowMinutes": number
    }
  ],
  "Description": "string",
  "Goal": {
    "AttainmentGoal": number,
    "Interval": { ... },
    "WarningThreshold": number
  },
  "RequestBasedSliConfig": {
    "ComparisonOperator": "string",
    "MetricThreshold": number,
    "RequestBasedSliMetricConfig": {
      "DependencyConfig": {
        "DependencyKeyAttributes": {
          "string" : "string"
        },
        "DependencyOperationName": "string"
      },
      "KeyAttributes": {
        "string" : "string"
      },
      "MetricType": "string",
      "MonitoredRequestCountMetric": { ... },
      "OperationName": "string",
      "TotalRequestCountMetric": [
        {
```

```

    "AccountId": "string",
    "Expression": "string",
    "Id": "string",
    "Label": "string",
    "MetricStat": {
      "Metric": {
        "Dimensions": [
          {
            "Name": "string",
            "Value": "string"
          }
        ],
        "MetricName": "string",
        "Namespace": "string"
      },
      "Period": number,
      "Stat": "string",
      "Unit": "string"
    },
    "Period": number,
    "ReturnData": boolean
  }
]
}
},
"SliConfig": {
  "ComparisonOperator": "string",
  "MetricThreshold": number,
  "SliMetricConfig": {
    "DependencyConfig": {
      "DependencyKeyAttributes": {
        "string" : "string"
      },
      "DependencyOperationName": "string"
    },
    "KeyAttributes": {
      "string" : "string"
    },
    "MetricDataQueries": [
      {
        "AccountId": "string",
        "Expression": "string",
        "Id": "string",
        "Label": "string",

```

```

    "MetricStat": {
      "Metric": {
        "Dimensions": [
          {
            "Name": "string",
            "Value": "string"
          }
        ],
        "MetricName": "string",
        "Namespace": "string"
      },
      "Period": number,
      "Stat": "string",
      "Unit": "string"
    },
    "Period": number,
    "ReturnData": boolean
  }
],
"MetricType": "string",
"OperationName": "string",
"PeriodSeconds": number,
"Statistic": "string"
}
}
}

```

URI Request Parameters

The request uses the following URI parameters.

Id

The Amazon Resource Name (ARN) or name of the service level objective that you want to update.

Pattern: `[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]$|
^arn:aws:application-signals:[^:]*:[^:]*:slo/[0-9A-Za-z][-._0-9A-Za-z]
{0,126}[0-9A-Za-z]`

Required: Yes

Request Body

The request accepts the following data in JSON format.

BurnRateConfigurations

Use this array to create *burn rates* for this SLO. Each burn rate is a metric that indicates how fast the service is consuming the error budget, relative to the attainment goal of the SLO.

Type: Array of [BurnRateConfiguration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Required: No

Description

An optional description for the SLO.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Goal

A structure that contains the attributes that determine the goal of the SLO. This includes the time period for evaluation and the attainment threshold.

Type: [Goal](#) object

Required: No

RequestBasedSliConfig

If this SLO is a request-based SLO, this structure defines the information about what performance metric this SLO will monitor.

You can't specify both `SliConfig` and `RequestBasedSliConfig` in the same operation.

Type: [RequestBasedServiceLevelIndicatorConfig](#) object

Required: No

SliConfig

If this SLO is a period-based SLO, this structure defines the information about what performance metric this SLO will monitor.

Type: [ServiceLevelIndicatorConfig](#) object

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "Slo": {
    "Arn": "string",
    "BurnRateConfigurations": [
      {
        "LookBackWindowMinutes": number
      }
    ],
    "CreatedTime": number,
    "Description": "string",
    "EvaluationType": "string",
    "Goal": {
      "AttainmentGoal": number,
      "Interval": { ... },
      "WarningThreshold": number
    },
    "LastUpdatedTime": number,
    "MetricSourceType": "string",
    "Name": "string",
    "RequestBasedSli": {
      "ComparisonOperator": "string",
      "MetricThreshold": number,
      "RequestBasedSliMetric": {
        "DependencyConfig": {
          "DependencyKeyAttributes": {
            "string" : "string"
          },
          "DependencyOperationName": "string"
        }
      }
    }
  }
}
```

```

    "KeyAttributes": {
      "string" : "string"
    },
    "MetricType": "string",
    "MonitoredRequestCountMetric": { ... },
    "OperationName": "string",
    "TotalRequestCountMetric": [
      {
        "AccountId": "string",
        "Expression": "string",
        "Id": "string",
        "Label": "string",
        "MetricStat": {
          "Metric": {
            "Dimensions": [
              {
                "Name": "string",
                "Value": "string"
              }
            ],
            "MetricName": "string",
            "Namespace": "string"
          },
          "Period": number,
          "Stat": "string",
          "Unit": "string"
        },
        "Period": number,
        "ReturnData": boolean
      }
    ]
  },
  "Sli": {
    "ComparisonOperator": "string",
    "MetricThreshold": number,
    "SliMetric": {
      "DependencyConfig": {
        "DependencyKeyAttributes": {
          "string" : "string"
        },
        "DependencyOperationName": "string"
      },
      "KeyAttributes": {

```



```

        "string" : "string"
    },
    "MetricDataQueries": [
        {
            "AccountId": "string",
            "Expression": "string",
            "Id": "string",
            "Label": "string",
            "MetricStat": {
                "Metric": {
                    "Dimensions": [
                        {
                            "Name": "string",
                            "Value": "string"
                        }
                    ],
                    "MetricName": "string",
                    "Namespace": "string"
                },
                "Period": number,
                "Stat": "string",
                "Unit": "string"
            },
            "Period": number,
            "ReturnData": boolean
        }
    ],
    "MetricType": "string",
    "OperationName": "string"
}
}
}
}

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Slo

A structure that contains information about the SLO that you just updated.

Type: [ServiceLevelObjective](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

ResourceNotFoundException

Resource not found.

HTTP Status Code: 404

ThrottlingException

The request was throttled because of quota limits.

HTTP Status Code: 429

ValidationException

The resource is not valid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

Data Types

The Amazon CloudWatch Application Signals API contains several data types that various actions use. This section describes each data type in detail.

Note

The order of each element in a data type structure is not guaranteed. Applications should not assume a particular order.

The following data types are supported:

- [BatchUpdateExclusionWindowsError](#)
- [BurnRateConfiguration](#)
- [CalendarInterval](#)
- [DependencyConfig](#)
- [Dimension](#)
- [ExclusionWindow](#)
- [Goal](#)
- [Interval](#)
- [Metric](#)
- [MetricDataQuery](#)
- [MetricReference](#)
- [MetricStat](#)
- [MonitoredRequestCountMetricDataQueries](#)
- [RecurrenceRule](#)
- [RequestBasedServiceLevelIndicator](#)
- [RequestBasedServiceLevelIndicatorConfig](#)
- [RequestBasedServiceLevelIndicatorMetric](#)
- [RequestBasedServiceLevelIndicatorMetricConfig](#)
- [RollingInterval](#)
- [Service](#)

- [ServiceDependency](#)
- [ServiceDependent](#)
- [ServiceLevelIndicator](#)
- [ServiceLevelIndicatorConfig](#)
- [ServiceLevelIndicatorMetric](#)
- [ServiceLevelIndicatorMetricConfig](#)
- [ServiceLevelObjective](#)
- [ServiceLevelObjectiveBudgetReport](#)
- [ServiceLevelObjectiveBudgetReportError](#)
- [ServiceLevelObjectiveSummary](#)
- [ServiceOperation](#)
- [ServiceSummary](#)
- [Tag](#)
- [Window](#)

BatchUpdateExclusionWindowsError

An array of structures, where each structure includes an error indicating that one of the requests in the array was not valid.

Contents

ErrorCode

The error code.

Type: String

Required: Yes

ErrorMessage

The error message.

Type: String

Required: Yes

SloId

The SLO ID in the error.

Type: String

Pattern: `[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]$|
^arn:aws:application-signals:[^:]*:[^:]*:slo/[0-9A-Za-z][-._0-9A-Za-z]
{0,126}[0-9A-Za-z]`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

BurnRateConfiguration

This object defines the length of the look-back window used to calculate one burn rate metric for this SLO. The burn rate measures how fast the service is consuming the error budget, relative to the attainment goal of the SLO. A burn rate of exactly 1 indicates that the SLO goal will be met exactly.

For example, if you specify 60 as the number of minutes in the look-back window, the burn rate is calculated as the following:

burn rate = error rate over the look-back window / (100% - attainment goal percentage)

For more information about burn rates, see [Calculate burn rates](#).

Contents

LookBackWindowMinutes

The number of minutes to use as the look-back window.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 10080.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CalendarInterval

If the interval for this service level objective is a calendar interval, this structure contains the interval specifications.

Contents

Duration

Specifies the duration of each calendar interval. For example, if `Duration` is 1 and `DurationUnit` is `MONTH`, each interval is one month, aligned with the calendar.

Type: Integer

Valid Range: Minimum value of 1.

Required: Yes

DurationUnit

Specifies the calendar interval unit.

Type: String

Valid Values: `MINUTE` | `HOUR` | `DAY` | `MONTH`

Required: Yes

StartTime

The date and time when you want the first interval to start. Be sure to choose a time that configures the intervals the way that you want. For example, if you want weekly intervals starting on Mondays at 6 a.m., be sure to specify a start time that is a Monday at 6 a.m.

When used in a raw HTTP Query API, it is formatted as be epoch time in seconds. For example: `1698778057`

As soon as one calendar interval ends, another automatically begins.

Type: Timestamp

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

DependencyConfig

Identifies the dependency using the `DependencyKeyAttributes` and `DependencyOperationName`.

When creating a service dependency SLO, you must specify the `KeyAttributes` of the service, and the `DependencyConfig` for the dependency. You can specify the `OperationName` of the service, from which it calls the dependency. Alternatively, you can exclude `OperationName` and the SLO will monitor all of the service's operations that call the dependency.

Contents

DependencyKeyAttributes

This is a string-to-string map. It can include the following fields.

- `Type` designates the type of object this is.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.
- `Identifier` identifies the resource objects of this resource. This is used only if the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Environment` specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: Yes

DependencyOperationName

The name of the called operation in the dependency.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Dimension

A dimension is a name/value pair that is part of the identity of a metric. Because dimensions are part of the unique identifier for a metric, whenever you add a unique name/value pair to one of your metrics, you are creating a new variation of that metric. For example, many Amazon EC2 metrics publish InstanceId as a dimension name, and the actual instance ID as the value for that dimension.

You can assign up to 30 dimensions to a metric.

Contents

Name

The name of the dimension. Dimension names must contain only ASCII characters, must include at least one non-whitespace character, and cannot start with a colon (:). ASCII control characters are not supported as part of dimension names.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

Value

The value of the dimension. Dimension values must contain only ASCII characters and must include at least one non-whitespace character. ASCII control characters are not supported as part of dimension values.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ExclusionWindow

The core SLO time window exclusion object that includes Window, StartTime, RecurrenceRule, and Reason.

Contents

Window

The SLO time window exclusion .

Type: [Window](#) object

Required: Yes

Reason

A description explaining why this time period should be excluded from SLO calculations.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

RecurrenceRule

The recurrence rule for the SLO time window exclusion. Supports both cron and rate expressions.

Type: [RecurrenceRule](#) object

Required: No

StartTime

The start of the SLO time window exclusion. Defaults to current time if not specified.

Type: Timestamp

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Goal

This structure contains the attributes that determine the goal of an SLO. This includes the time period for evaluation and the attainment threshold.

Contents

AttainmentGoal

The threshold that determines if the goal is being met.

If this is a period-based SLO, the attainment goal is the percentage of good periods that meet the threshold requirements to the total periods within the interval. For example, an attainment goal of 99.9% means that within your interval, you are targeting 99.9% of the periods to be in healthy state.

If this is a request-based SLO, the attainment goal is the percentage of requests that must be successful to meet the attainment goal.

If you omit this parameter, 99 is used to represent 99% as the attainment goal.

Type: Double

Required: No

Interval

The time period used to evaluate the SLO. It can be either a calendar interval or rolling interval.

If you omit this parameter, a rolling interval of 7 days is used.

Type: [Interval](#) object

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: No

WarningThreshold

The percentage of remaining budget over total budget that you want to get warnings for. If you omit this parameter, the default of 50.0 is used.

Type: Double

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Interval

The time period used to evaluate the SLO. It can be either a calendar interval or rolling interval.

Contents

Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

CalendarInterval

If the interval is a calendar interval, this structure contains the interval specifications.

Type: [CalendarInterval](#) object

Required: No

RollingInterval

If the interval is a rolling interval, this structure contains the interval specifications.

Type: [RollingInterval](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Metric

This structure defines the metric used for a service level indicator, including the metric name, namespace, and dimensions

Contents

Dimensions

An array of one or more dimensions to use to define the metric that you want to use. For more information, see [Dimensions](#).

Type: Array of [Dimension](#) objects

Array Members: Minimum number of 0 items. Maximum number of 30 items.

Required: No

MetricName

The name of the metric to use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

Namespace

The namespace of the metric. For more information, see [Namespaces](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `. * [^ :] . *`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetricDataQuery

Use this structure to define a metric or metric math expression that you want to use as for a service level objective.

Each `MetricDataQuery` in the `MetricDataQueries` array specifies either a metric to retrieve, or a metric math expression to be performed on retrieved metrics. A single `MetricDataQueries` array can include as many as 20 `MetricDataQuery` structures in the array. The 20 structures can include as many as 10 structures that contain a `MetricStat` parameter to retrieve a metric, and as many as 10 structures that contain the `Expression` parameter to perform a math expression. Of those `Expression` structures, exactly one must have `true` as the value for `ReturnData`. The result of this expression used for the SLO.

For more information about metric math expressions, see [CloudWatchUse metric math](#).

Within each `MetricDataQuery` object, you must specify either `Expression` or `MetricStat` but not both.

Contents

Id

A short name used to tie this object to the results in the response. This `Id` must be unique within a `MetricDataQueries` array. If you are performing math expressions on this set of data, this name represents that data and can serve as a variable in the metric math expression. The valid characters are letters, numbers, and underscore. The first character must be a lowercase letter.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

AccountId

The ID of the account where this metric is located. If you are performing this operation in a monitoring account, use this to specify which source account to retrieve this metric from.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

Expression

This field can contain a metric math expression to be performed on the other metrics that you are retrieving within this `MetricDataQueries` structure.

A math expression can use the `Id` of the other metrics or queries to refer to those metrics, and can also use the `Id` of other expressions to use the result of those expressions. For more information about metric math expressions, see [Metric Math Syntax and Functions](#) in the *Amazon CloudWatch User Guide*.

Within each `MetricDataQuery` object, you must specify either `Expression` or `MetricStat` but not both.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

Label

A human-readable label for this metric or expression. This is especially useful if this is an expression, so that you know what the value represents. If the metric or expression is shown in a CloudWatch dashboard widget, the label is shown. If `Label` is omitted, CloudWatch generates a default.

You can put dynamic expressions into a label, so that it is more descriptive. For more information, see [Using Dynamic Labels](#).

Type: String

Required: No

MetricStat

A metric to be used directly for the SLO, or to be used in the math expression that will be used for the SLO.

Within one `MetricDataQuery` object, you must specify either `Expression` or `MetricStat` but not both.

Type: [MetricStat](#) object

Required: No

Period

The granularity, in seconds, of the returned data points for this metric. For metrics with regular resolution, a period can be as short as one minute (60 seconds) and must be a multiple of 60. For high-resolution metrics that are collected at intervals of less than one minute, the period can be 1, 5, 10, 30, 60, or any multiple of 60. High-resolution metrics are those metrics stored by a `PutMetricData` call that includes a `StorageResolution` of 1 second.

If the `StartTime` parameter specifies a time stamp that is greater than 3 hours ago, you must specify the period as follows or no data points in that time range is returned:

- Start time between 3 hours and 15 days ago - Use a multiple of 60 seconds (1 minute).
- Start time between 15 and 63 days ago - Use a multiple of 300 seconds (5 minutes).
- Start time greater than 63 days ago - Use a multiple of 3600 seconds (1 hour).

Type: Integer

Valid Range: Minimum value of 1.

Required: No

ReturnData

Use this only if you are using a metric math expression for the SLO. Specify `true` for `ReturnData` for only the one expression result to use as the alarm. For all other metrics and expressions in the same `CreateServiceLevelObjective` operation, specify `ReturnData` as `false`.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetricReference

This structure contains information about one CloudWatch metric associated with this entity discovered by Application Signals.

Contents

MetricName

The name of the metric.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

MetricType

Used to display the appropriate statistics in the CloudWatch console.

Type: String

Pattern: [A-Za-z0-9 -]+

Required: Yes

Namespace

The namespace of the metric. For more information, see [CloudWatchNamespaces](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: .*[^\:]*

Required: Yes

AccountId

AWS account ID.

Type: String

Pattern: [0-9]{12}

Required: No

Dimensions

An array of one or more dimensions that further define the metric. For more information, see [CloudWatchDimensions](#).

Type: Array of [Dimension](#) objects

Array Members: Minimum number of 0 items. Maximum number of 30 items.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetricStat

This structure defines the metric to be used as the service level indicator, along with the statistics, period, and unit.

Contents

Metric

The metric to use as the service level indicator, including the metric name, namespace, and dimensions.

Type: [Metric](#) object

Required: Yes

Period

The granularity, in seconds, to be used for the metric. For metrics with regular resolution, a period can be as short as one minute (60 seconds) and must be a multiple of 60. For high-resolution metrics that are collected at intervals of less than one minute, the period can be 1, 5, 10, 30, 60, or any multiple of 60. High-resolution metrics are those metrics stored by a PutMetricData call that includes a StorageResolution of 1 second.

Type: Integer

Valid Range: Minimum value of 1.

Required: Yes

Stat

The statistic to use for comparison to the threshold. It can be any CloudWatch statistic or extended statistic. For more information about statistics, see [CloudWatch statistics definitions](#).

Type: String

Required: Yes

Unit

If you omit Unit then all data that was collected with any unit is returned, along with the corresponding units that were specified when the data was reported to CloudWatch. If you

specify a unit, the operation returns only data that was collected with that unit specified. If you specify a unit that does not match the data collected, the results of the operation are null. CloudWatch does not perform unit conversions.

Type: String

Valid Values: Microseconds | Milliseconds | Seconds | Bytes | Kilobytes | Megabytes | Gigabytes | Terabytes | Bits | Kilobits | Megabits | Gigabits | Terabits | Percent | Count | Bytes/Second | Kilobytes/Second | Megabytes/Second | Gigabytes/Second | Terabytes/Second | Bits/Second | Kilobits/Second | Megabits/Second | Gigabits/Second | Terabits/Second | Count/Second | None

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MonitoredRequestCountMetricDataQueries

This structure defines the metric that is used as the "good request" or "bad request" value for a request-based SLO. This value observed for the metric defined in `TotalRequestCountMetric` is divided by the number found for `MonitoredRequestCountMetric` to determine the percentage of successful requests that this SLO tracks.

Contents

Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

BadCountMetric

If you want to count "bad requests" to determine the percentage of successful requests for this request-based SLO, specify the metric to use as "bad requests" in this structure.

Type: Array of [MetricDataQuery](#) objects

Required: No

GoodCountMetric

If you want to count "good requests" to determine the percentage of successful requests for this request-based SLO, specify the metric to use as "good requests" in this structure.

Type: Array of [MetricDataQuery](#) objects

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

RecurrenceRule

The recurrence rule for the SLO time window exclusion .

Contents

Expression

A cron or rate expression that specifies the schedule for the exclusion window.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RequestBasedServiceLevelIndicator

This structure contains information about the performance metric that a request-based SLO monitors.

Contents

RequestBasedSliMetric

A structure that contains information about the metric that the SLO monitors.

Type: [RequestBasedServiceLevelIndicatorMetric](#) object

Required: Yes

ComparisonOperator

The arithmetic operation used when comparing the specified metric to the threshold.

Type: String

Valid Values: `GreaterThanOrEqualTo` | `GreaterThan` | `LessThan` | `LessThanOrEqualTo`

Required: No

MetricThreshold

This value is the threshold that the observed metric values of the SLI metric are compared to.

Type: Double

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RequestBasedServiceLevelIndicatorConfig

This structure specifies the information about the service and the performance metric that a request-based SLO is to monitor.

Contents

RequestBasedSliMetricConfig

Use this structure to specify the metric to be used for the SLO.

Type: [RequestBasedServiceLevelIndicatorMetricConfig](#) object

Required: Yes

ComparisonOperator

The arithmetic operation to use when comparing the specified metric to the threshold. This parameter is required if this SLO is tracking the Latency metric.

Type: String

Valid Values: `GreaterThanOrEqualTo` | `GreaterThan` | `LessThan` | `LessThanOrEqualTo`

Required: No

MetricThreshold

The value that the SLI metric is compared to. This parameter is required if this SLO is tracking the Latency metric.

Type: Double

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RequestBasedServiceLevelIndicatorMetric

This structure contains the information about the metric that is used for a request-based SLO.

Contents

MonitoredRequestCountMetric

This structure defines the metric that is used as the "good request" or "bad request" value for a request-based SLO. This value observed for the metric defined in `TotalRequestCountMetric` is divided by the number found for `MonitoredRequestCountMetric` to determine the percentage of successful requests that this SLO tracks.

Type: [MonitoredRequestCountMetricDataQueries](#) object

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: Yes

TotalRequestCountMetric

This structure defines the metric that is used as the "total requests" number for a request-based SLO. The number observed for this metric is divided by the number of "good requests" or "bad requests" that is observed for the metric defined in `MonitoredRequestCountMetric`.

Type: Array of [MetricDataQuery](#) objects

Required: Yes

DependencyConfig

Identifies the dependency using the `DependencyKeyAttributes` and `DependencyOperationName`.

Type: [DependencyConfig](#) object

Required: No

KeyAttributes

This is a string-to-string map that contains information about the type of object that this SLO is related to. It can include the following fields.

- `Type` designates the type of object that this SLO is related to.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.
- `Identifier` identifies the resource objects of this resource. This is used only if the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Environment` specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: No

MetricType

If the SLO monitors either the LATENCY or AVAILABILITY metric that Application Signals collects, this field displays which of those metrics is used.

Type: String

Valid Values: LATENCY | AVAILABILITY

Required: No

OperationName

If the SLO monitors a specific operation of the service, this field displays that operation name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RequestBasedServiceLevelIndicatorMetricConfig

Use this structure to specify the information for the metric that a period-based SLO will monitor.

Contents

DependencyConfig

Identifies the dependency using the `DependencyKeyAttributes` and `DependencyOperationName`.

Type: [DependencyConfig](#) object

Required: No

KeyAttributes

If this SLO is related to a metric collected by Application Signals, you must use this field to specify which service the SLO metric is related to. To do so, you must specify at least the `Type`, `Name`, and `Environment` attributes.

This is a string-to-string map. It can include the following fields.

- `Type` designates the type of object this is.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.
- `Identifier` identifies the resource objects of this resource. This is used only if the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Environment` specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: No

MetricType

If the SLO is to monitor either the LATENCY or AVAILABILITY metric that Application Signals collects, use this field to specify which of those metrics is used.

Type: String

Valid Values: LATENCY | AVAILABILITY

Required: No

MonitoredRequestCountMetric

Use this structure to define the metric that you want to use as the "good request" or "bad request" value for a request-based SLO. This value observed for the metric defined in `TotalRequestCountMetric` will be divided by the number found for `MonitoredRequestCountMetric` to determine the percentage of successful requests that this SLO tracks.

Type: [MonitoredRequestCountMetricDataQueries](#) object

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: No

OperationName

If the SLO is to monitor a specific operation of the service, use this field to specify the name of that operation.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

TotalRequestCountMetric

Use this structure to define the metric that you want to use as the "total requests" number for a request-based SLO. This result will be divided by the "good request" or "bad request" value defined in `MonitoredRequestCountMetric`.

Type: Array of [MetricDataQuery](#) objects

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RollingInterval

If the interval for this SLO is a rolling interval, this structure contains the interval specifications.

Contents

Duration

Specifies the duration of each rolling interval. For example, if `Duration` is 7 and `DurationUnit` is DAY, each rolling interval is seven days.

Type: Integer

Valid Range: Minimum value of 1.

Required: Yes

DurationUnit

Specifies the rolling interval unit.

Type: String

Valid Values: MINUTE | HOUR | DAY | MONTH

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Service

This structure contains information about one of your services that was discovered by Application Signals.

Contents

KeyAttributes

This is a string-to-string map. It can include the following fields.

- `Type` designates the type of object this is.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.
- `Identifier` identifies the resource objects of this resource. This is used only if the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Environment` specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: Yes

MetricReferences

An array of structures that each contain information about one metric associated with this service.

Type: Array of [MetricReference](#) objects

Required: Yes

AttributeMaps

This structure contains one or more string-to-string maps that help identify this service. It can include *platform attributes*, *application attributes*, and *telemetry attributes*.

Platform attributes contain information the service's platform.

- `PlatformType` defines the hosted-in platform.
- `EKS.Cluster` is the name of the Amazon EKS cluster.
- `K8s.Cluster` is the name of the self-hosted Kubernetes cluster.
- `K8s.Namespace` is the name of the Kubernetes namespace in either Amazon EKS or Kubernetes clusters.
- `K8s.Workload` is the name of the Kubernetes workload in either Amazon EKS or Kubernetes clusters.
- `K8s.Node` is the name of the Kubernetes node in either Amazon EKS or Kubernetes clusters.
- `K8s.Pod` is the name of the Kubernetes pod in either Amazon EKS or Kubernetes clusters.
- `EC2.AutoScalingGroup` is the name of the Amazon EC2 Auto Scaling group.
- `EC2.InstanceId` is the ID of the Amazon EC2 instance.
- `Host` is the name of the host, for all platform types.

Application attributes contain information about the application.

- `AWS.Application` is the application's name in AWS Service Catalog AppRegistry.
- `AWS.Application.ARN` is the application's ARN in AWS Service Catalog AppRegistry.

Telemetry attributes contain telemetry information.

- `Telemetry.SDK` is the fingerprint of the OpenTelemetry SDK version for instrumented services.
- `Telemetry.Agent` is the fingerprint of the agent used to collect and send telemetry data.
- `Telemetry.Source` Specifies the point of application where the telemetry was collected or specifies what was used for the source of telemetry data.

Type: Array of string to string maps

Required: No

LogGroupReferences

An array of string-to-string maps that each contain information about one log group associated with this service. Each string-to-string map includes the following fields:

- "Type": "AWS::Resource"
- "ResourceType": "AWS::Logs::LogGroup"
- "Identifier": "*name-of-log-group*"

Type: Array of string to string maps

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!~]+[-~]*`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServiceDependency

This structure contains information about one dependency of this service.

Contents

DependencyKeyAttributes

This is a string-to-string map. It can include the following fields.

- `Type` designates the type of object this is.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.
- `Identifier` identifies the resource objects of this resource. This is used only if the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Environment` specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: Yes

DependencyOperationName

The name of the called operation in the dependency.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

MetricReferences

An array of structures that each contain information about one metric associated with this service dependency that was discovered by Application Signals.

Type: Array of [MetricReference](#) objects

Required: Yes

OperationName

The name of the operation in this service that calls the dependency.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServiceDependent

This structure contains information about a service dependent that was discovered by Application Signals. A dependent is an entity that invoked the specified service during the provided time range. Dependents include other services, CloudWatch Synthetics canaries, and clients that are instrumented with CloudWatch RUM app monitors.

Contents

DependentKeyAttributes

This is a string-to-string map. It can include the following fields.

- `Type` designates the type of object this is.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.
- `Identifier` identifies the resource objects of this resource. This is used only if the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Environment` specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: Yes

MetricReferences

An array of structures that each contain information about one metric associated with this service dependent that was discovered by Application Signals.

Type: Array of [MetricReference](#) objects

Required: Yes

DependentOperationName

If the dependent invoker was a service that invoked it from an operation, the name of that dependent operation is displayed here.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

OperationName

If the invoked entity is an operation on an entity, the name of that dependent operation is displayed here.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServiceLevelIndicator

This structure contains information about the performance metric that a period-based SLO monitors.

Contents

ComparisonOperator

The arithmetic operation used when comparing the specified metric to the threshold.

Type: String

Valid Values: `GreaterThanOrEqualTo` | `GreaterThan` | `LessThan` | `LessThanOrEqualTo`

Required: Yes

MetricThreshold

The value that the SLI metric is compared to.

Type: Double

Required: Yes

SLIMetric

A structure that contains information about the metric that the SLO monitors.

Type: [ServiceLevelIndicatorMetric](#) object

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServiceLevelIndicatorConfig

This structure specifies the information about the service and the performance metric that a period-based SLO is to monitor.

Contents

ComparisonOperator

The arithmetic operation to use when comparing the specified metric to the threshold.

Type: String

Valid Values: `GreaterThanOrEqualTo` | `GreaterThan` | `LessThan` | `LessThanOrEqualTo`

Required: Yes

MetricThreshold

This parameter is used only when a request-based SLO tracks the Latency metric. Specify the threshold value that the observed Latency metric values are to be compared to.

Type: Double

Required: Yes

SliMetricConfig

Use this structure to specify the metric to be used for the SLO.

Type: [ServiceLevelIndicatorMetricConfig](#) object

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

ServiceLevelIndicatorMetric

This structure contains the information about the metric that is used for a period-based SLO.

Contents

MetricDataQueries

If this SLO monitors a CloudWatch metric or the result of a CloudWatch metric math expression, this structure includes the information about that metric or expression.

Type: Array of [MetricDataQuery](#) objects

Required: Yes

DependencyConfig

Identifies the dependency using the `DependencyKeyAttributes` and `DependencyOperationName`.

Type: [DependencyConfig](#) object

Required: No

KeyAttributes

This is a string-to-string map that contains information about the type of object that this SLO is related to. It can include the following fields.

- `Type` designates the type of object that this SLO is related to.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.
- `Identifier` identifies the resource objects of this resource. This is used only if the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Environment` specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: No

MetricType

If the SLO monitors either the LATENCY or AVAILABILITY metric that Application Signals collects, this field displays which of those metrics is used.

Type: String

Valid Values: LATENCY | AVAILABILITY

Required: No

OperationName

If the SLO monitors a specific operation of the service, this field displays that operation name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServiceLevelIndicatorMetricConfig

Use this structure to specify the information for the metric that a period-based SLO will monitor.

Contents

DependencyConfig

Identifies the dependency using the `DependencyKeyAttributes` and `DependencyOperationName`.

Type: [DependencyConfig](#) object

Required: No

KeyAttributes

If this SLO is related to a metric collected by Application Signals, you must use this field to specify which service the SLO metric is related to. To do so, you must specify at least the `Type`, `Name`, and `Environment` attributes.

This is a string-to-string map. It can include the following fields.

- `Type` designates the type of object this is.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.
- `Identifier` identifies the resource objects of this resource. This is used only if the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Environment` specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: No

MetricDataQueries

If this SLO monitors a CloudWatch metric or the result of a CloudWatch metric math expression, use this structure to specify that metric or expression.

Type: Array of [MetricDataQuery](#) objects

Required: No

MetricType

If the SLO is to monitor either the LATENCY or AVAILABILITY metric that Application Signals collects, use this field to specify which of those metrics is used.

Type: String

Valid Values: LATENCY | AVAILABILITY

Required: No

OperationName

If the SLO is to monitor a specific operation of the service, use this field to specify the name of that operation.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

PeriodSeconds

The number of seconds to use as the period for SLO evaluation. Your application's performance is compared to the SLI during each period. For each period, the application is determined to have either achieved or not achieved the necessary performance.

Type: Integer

Valid Range: Minimum value of 60. Maximum value of 900.

Required: No

Statistic

The statistic to use for comparison to the threshold. It can be any CloudWatch statistic or extended statistic. For more information about statistics, see [CloudWatch statistics definitions](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 20.

Pattern: [a-zA-Z0-9.]+

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServiceLevelObjective

A structure containing information about one service level objective (SLO) that has been created in Application Signals. Creating SLOs can help you ensure your services are performing to the level that you expect. SLOs help you set and track a specific target level for the reliability and availability of your applications and services. Each SLO uses a service level indicator (SLI), which is a key performance metric, to calculate how much underperformance can be tolerated before the goal that you set for the SLO is not achieved.

Contents

Arn

The ARN of this SLO.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:aws:application-signals:[^:]*:[^:]*:slo/[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]`

Required: Yes

CreatedTime

The date and time that this SLO was created. When used in a raw HTTP Query API, it is formatted as `yyyy-MM-dd'T'HH:mm:ss`. For example, `2019-07-01T23:59:59`.

Type: Timestamp

Required: Yes

Goal

This structure contains the attributes that determine the goal of an SLO. This includes the time period for evaluation and the attainment threshold.

Type: [Goal](#) object

Required: Yes

LastUpdatedTime

The time that this SLO was most recently updated. When used in a raw HTTP Query API, it is formatted as yyyy-MM-dd 'T' HH:mm:ss. For example, 2019-07-01T23:59:59.

Type: Timestamp

Required: Yes

Name

The name of this SLO.

Type: String

Pattern: [0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]

Required: Yes

BurnRateConfigurations

Each object in this array defines the length of the look-back window used to calculate one burn rate metric for this SLO. The burn rate measures how fast the service is consuming the error budget, relative to the attainment goal of the SLO.

Type: Array of [BurnRateConfiguration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Required: No

Description

The description that you created for this SLO.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

EvaluationType

Displays whether this is a period-based SLO or a request-based SLO.

Type: String

Valid Values: PeriodBased | RequestBased

Required: No

MetricSourceType

Displays the SLI metric source type for this SLO. Supported types are:

- Service operation
- Service dependency
- CloudWatch metric

Type: String

Valid Values: ServiceOperation | CloudWatchMetric | ServiceDependency

Required: No

RequestBasedSli

A structure containing information about the performance metric that this SLO monitors, if this is a request-based SLO.

Type: [RequestBasedServiceLevelIndicator](#) object

Required: No

Sli

A structure containing information about the performance metric that this SLO monitors, if this is a period-based SLO.

Type: [ServiceLevelIndicator](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

ServiceLevelObjectiveBudgetReport

A structure containing an SLO budget report that you have requested.

Contents

Arn

The ARN of the SLO that this report is for.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:aws:application-signals:[^:]*:[^:]*:slo/[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]`

Required: Yes

BudgetStatus

The status of this SLO, as it relates to the error budget for the entire time interval.

- OK means that the SLO had remaining budget above the warning threshold, as of the time that you specified in `TimeStamp`.
- WARNING means that the SLO's remaining budget was below the warning threshold, as of the time that you specified in `TimeStamp`.
- BREACHED means that the SLO's budget was exhausted, as of the time that you specified in `TimeStamp`.
- INSUFFICIENT_DATA means that the specified start and end times were before the SLO was created, or that attainment data is missing.

Type: String

Valid Values: OK | WARNING | BREACHED | INSUFFICIENT_DATA

Required: Yes

Name

The name of the SLO that this report is for.

Type: String

Pattern: [0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]

Required: Yes

Attainment

A number between 0 and 100 that represents the success percentage of your application compared to the goal set by the SLO.

If this is a period-based SLO, the number is the percentage of time periods that the service has attained the SLO's attainment goal, as of the time of the request.

If this is a request-based SLO, the number is the number of successful requests divided by the number of total requests, multiplied by 100, during the time range that you specified in your request.

Type: Double

Required: No

BudgetRequestsRemaining

This field is displayed only for request-based SLOs. It displays the number of failed requests that can be tolerated before any more successful requests occur, and still have the application meet its SLO goal.

This number can go up and down between different reports, based on both how many successful requests and how many failed requests occur in that time.

Type: Integer

Required: No

BudgetSecondsRemaining

The budget amount remaining before the SLO status becomes BREACHING, at the time specified in the `Timestamp` parameter of the request. If this value is negative, then the SLO is already in BREACHING status.

This field is included only if the SLO is a period-based SLO.

Type: Integer

Required: No

EvaluationType

Displays whether this budget report is for a period-based SLO or a request-based SLO.

Type: String

Valid Values: PeriodBased | RequestBased

Required: No

Goal

This structure contains the attributes that determine the goal of an SLO. This includes the time period for evaluation and the attainment threshold.

Type: [Goal](#) object

Required: No

RequestBasedSli

This structure contains information about the performance metric that a request-based SLO monitors.

Type: [RequestBasedServiceLevelIndicator](#) object

Required: No

Sli

A structure that contains information about the performance metric that this SLO monitors.

Type: [ServiceLevelIndicator](#) object

Required: No

TotalBudgetRequests

This field is displayed only for request-based SLOs. It displays the total number of failed requests that can be tolerated during the time range between the start of the interval and the time stamp supplied in the budget report request. It is based on the total number of requests that occurred, and the percentage specified in the attainment goal. If the number of failed requests matches this number or is higher, then this SLO is currently breaching.

This number can go up and down between reports with different time stamps, based on both how many total requests occur.

Type: Integer

Required: No

TotalBudgetSeconds

The total number of seconds in the error budget for the interval. This field is included only if the SLO is a period-based SLO.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServiceLevelObjectiveBudgetReportError

A structure containing information about one error that occurred during a [BatchGetServiceLevelObjectiveBudgetReport](#) operation.

Contents

Arn

The ARN of the SLO that this error is related to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:aws:application-signals:[^:]*:[^:]*:slo/[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]`

Required: Yes

ErrorCode

The error code for this error.

Type: String

Required: Yes

ErrorMessage

The message for this error.

Type: String

Required: Yes

Name

The name of the SLO that this error is related to.

Type: String

Pattern: `[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServiceLevelObjectiveSummary

A structure that contains information about one service level objective (SLO) created in Application Signals.

Contents

Arn

The ARN of this service level objective.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:aws:application-signals:[^:]*:[^:]*:slo/[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]`

Required: Yes

Name

The name of the service level objective.

Type: String

Pattern: `[0-9A-Za-z][-._0-9A-Za-z]{0,126}[0-9A-Za-z]`

Required: Yes

CreatedTime

The date and time that this service level objective was created. It is expressed as the number of milliseconds since Jan 1, 1970 00:00:00 UTC.

Type: Timestamp

Required: No

DependencyConfig

Identifies the dependency using the `DependencyKeyAttributes` and `DependencyOperationName`.

Type: [DependencyConfig](#) object

Required: No

EvaluationType

Displays whether this is a period-based SLO or a request-based SLO.

Type: String

Valid Values: PeriodBased | RequestBased

Required: No

KeyAttributes

This is a string-to-string map. It can include the following fields.

- Type designates the type of object this service level objective is for.
- ResourceType specifies the type of the resource. This field is used only when the value of the Type field is Resource or AWS::Resource.
- Name specifies the name of the object. This is used only if the value of the Type field is Service, RemoteService, or AWS::Service.
- Identifier identifies the resource objects of this resource. This is used only if the value of the Type field is Resource or AWS::Resource.
- Environment specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: [a-zA-Z]{1,50}

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: [-~]*[!-~]+[-~]*

Required: No

MetricSourceType

Displays the SLI metric source type for this SLO. Supported types are:

- Service operation

- Service dependency
- CloudWatch metric

Type: String

Valid Values: ServiceOperation | CloudWatchMetric | ServiceDependency

Required: No

OperationName

If this service level objective is specific to a single operation, this field displays the name of that operation.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServiceOperation

This structure contains information about an operation discovered by Application Signals. An operation is a specific function performed by a service that was discovered by Application Signals, and is often an API that is called by an upstream dependent.

Contents

MetricReferences

An array of structures that each contain information about one metric associated with this service operation that was discovered by Application Signals.

Type: Array of [MetricReference](#) objects

Required: Yes

Name

The name of the operation, discovered by Application Signals.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServiceSummary

This structure contains information about one of your services that was discovered by Application Signals

Contents

KeyAttributes

This is a string-to-string map that help identify the objects discovered by Application Signals. It can include the following fields.

- `Type` designates the type of object this is.
- `ResourceType` specifies the type of the resource. This field is used only when the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Name` specifies the name of the object. This is used only if the value of the `Type` field is `Service`, `RemoteService`, or `AWS::Service`.
- `Identifier` identifies the resource objects of this resource. This is used only if the value of the `Type` field is `Resource` or `AWS::Resource`.
- `Environment` specifies the location where this object is hosted, or what it belongs to.

Type: String to string map

Map Entries: Maximum number of 4 items.

Key Pattern: `[a-zA-Z]{1,50}`

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: `[-~]*[!-~]+[-~]*`

Required: Yes

MetricReferences

An array of structures that each contain information about one metric associated with this service.

Type: Array of [MetricReference](#) objects

Required: Yes

AttributeMaps

This structure contains one or more string-to-string maps that help identify this service. It can include *platform attributes*, *application attributes*, and *telemetry attributes*.

Platform attributes contain information the service's platform.

- `PlatformType` defines the hosted-in platform.
- `EKS.Cluster` is the name of the Amazon EKS cluster.
- `K8s.Cluster` is the name of the self-hosted Kubernetes cluster.
- `K8s.Namespace` is the name of the Kubernetes namespace in either Amazon EKS or Kubernetes clusters.
- `K8s.Workload` is the name of the Kubernetes workload in either Amazon EKS or Kubernetes clusters.
- `K8s.Node` is the name of the Kubernetes node in either Amazon EKS or Kubernetes clusters.
- `K8s.Pod` is the name of the Kubernetes pod in either Amazon EKS or Kubernetes clusters.
- `EC2.AutoScalingGroup` is the name of the Amazon EC2 Auto Scaling group.
- `EC2.InstanceId` is the ID of the Amazon EC2 instance.
- `Host` is the name of the host, for all platform types.

Application attributes contain information about the application.

- `AWS.Application` is the application's name in AWS Service Catalog AppRegistry.
- `AWS.Application.ARN` is the application's ARN in AWS Service Catalog AppRegistry.

Telemetry attributes contain telemetry information.

- `Telemetry.SDK` is the fingerprint of the OpenTelemetry SDK version for instrumented services.
- `Telemetry.Agent` is the fingerprint of the agent used to collect and send telemetry data.
- `Telemetry.Source` Specifies the point of application where the telemetry was collected or specifies what was used for the source of telemetry data.

Type: Array of string to string maps

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Tag

A key-value pair associated with a resource. Tags can help you organize and categorize your resources.

Contents

Key

A string that you can use to assign a value. The combination of tag keys and values can help you organize and categorize your resources.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Required: Yes

Value

The value for the specified tag key.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Window

The object that defines the time length of an exclusion window.

Contents

Duration

The number of time units for the exclusion window length.

Type: Integer

Valid Range: Minimum value of 1.

Required: Yes

DurationUnit

The unit of time for the exclusion window duration. Valid values: MINUTE, HOUR, DAY, MONTH.

Type: String

Valid Values: MINUTE | HOUR | DAY | MONTH

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Common Parameters

The following list contains the parameters that all actions use for signing Signature Version 4 requests with a query string. Any action-specific parameters are listed in the topic for that action. For more information about Signature Version 4, see [Signing AWS API requests](#) in the *IAM User Guide*.

Action

The action to be performed.

Type: string

Required: Yes

Version

The API version that the request is written for, expressed in the format YYYY-MM-DD.

Type: string

Required: Yes

X-Amz-Algorithm

The hash algorithm that you used to create the request signature.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Valid Values: AWS4-HMAC-SHA256

Required: Conditional

X-Amz-Credential

The credential scope value, which is a string that includes your access key, the date, the region you are targeting, the service you are requesting, and a termination string ("aws4_request"). The value is expressed in the following format: *access_key/YYYYMMDD/region/service/aws4_request*.

For more information, see [Create a signed AWS API request](#) in the *IAM User Guide*.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

X-Amz-Date

The date that is used to create the signature. The format must be ISO 8601 basic format (YYYYMMDD'T'HHMMSS'Z'). For example, the following date time is a valid X-Amz-Date value: 20120325T120000Z.

Condition: X-Amz-Date is optional for all requests; it can be used to override the date used for signing requests. If the Date header is specified in the ISO 8601 basic format, X-Amz-Date is not required. When X-Amz-Date is used, it always overrides the value of the Date header. For more information, see [Elements of an AWS API request signature](#) in the *IAM User Guide*.

Type: string

Required: Conditional

X-Amz-Security-Token

The temporary security token that was obtained through a call to AWS Security Token Service (AWS STS). For a list of services that support temporary security credentials from AWS STS, see [AWS services that work with IAM](#) in the *IAM User Guide*.

Condition: If you're using temporary security credentials from AWS STS, you must include the security token.

Type: string

Required: Conditional

X-Amz-Signature

Specifies the hex-encoded signature that was calculated from the string to sign and the derived signing key.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

X-Amz-SignedHeaders

Specifies all the HTTP headers that were included as part of the canonical request. For more information about specifying signed headers, see [Create a signed AWS API request](#) in the *IAM User Guide*.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

Common Errors

This section lists the errors common to the API actions of all AWS services. For errors specific to an API action for this service, see the topic for that API action.

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 403

ExpiredTokenException

The security token included in the request is expired

HTTP Status Code: 403

IncompleteSignature

The request signature does not conform to AWS standards.

HTTP Status Code: 403

InternalFailure

The request processing has failed because of an unknown error, exception or failure.

HTTP Status Code: 500

MalformedHttpRequestException

Problems with the request at the HTTP level, e.g. we can't decompress the body according to the decompression algorithm specified by the content-encoding.

HTTP Status Code: 400

NotAuthorized

You do not have permission to perform this action.

HTTP Status Code: 401

OptInRequired

The AWS access key ID needs a subscription for the service.

HTTP Status Code: 403

RequestAbortedException

Convenient exception that can be used when a request is aborted before a reply is sent back (e.g. client closed connection).

HTTP Status Code: 400

RequestEntityTooLargeException

Problems with the request at the HTTP level. The request entity is too large.

HTTP Status Code: 413

RequestExpired

The request reached the service more than 15 minutes after the date stamp on the request or more than 15 minutes after the request expiration date (such as for pre-signed URLs), or the date stamp on the request is more than 15 minutes in the future.

HTTP Status Code: 400

RequestTimeoutException

Problems with the request at the HTTP level. Reading the Request timed out.

HTTP Status Code: 408

ServiceUnavailable

The request has failed due to a temporary failure of the server.

HTTP Status Code: 503

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

UnrecognizedClientException

The X.509 certificate or AWS access key ID provided does not exist in our records.

HTTP Status Code: 403

UnknownOperationException

The action or operation requested is invalid. Verify that the action is typed correctly.

HTTP Status Code: 404

ValidationError

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400