



API Reference

AWS X-Ray



API Version 2016-04-12

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

AWS X-Ray: API Reference

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	1
Actions	2
BatchGetTraces	4
Request Syntax	4
URI Request Parameters	4
Request Body	4
Response Syntax	5
Response Elements	5
Errors	6
See Also	6
CancelTraceRetrieval	8
Request Syntax	8
URI Request Parameters	8
Request Body	8
Response Syntax	8
Response Elements	8
Errors	9
See Also	9
CreateGroup	10
Request Syntax	10
URI Request Parameters	10
Request Body	10
Response Syntax	12
Response Elements	12
Errors	12
See Also	13
CreateSamplingRule	14
Request Syntax	14
URI Request Parameters	15
Request Body	15
Response Syntax	15
Response Elements	16
Errors	16
See Also	17

DeleteGroup	18
Request Syntax	18
URI Request Parameters	18
Request Body	18
Response Syntax	19
Response Elements	19
Errors	19
See Also	19
DeleteResourcePolicy	21
Request Syntax	21
URI Request Parameters	21
Request Body	21
Response Syntax	22
Response Elements	22
Errors	22
See Also	22
DeleteSamplingRule	24
Request Syntax	24
URI Request Parameters	24
Request Body	24
Response Syntax	24
Response Elements	25
Errors	25
See Also	26
GetEncryptionConfig	27
Request Syntax	27
URI Request Parameters	27
Request Body	27
Response Syntax	27
Response Elements	27
Errors	28
See Also	28
GetGroup	29
Request Syntax	29
URI Request Parameters	29
Request Body	29

Response Syntax	30
Response Elements	30
Errors	30
See Also	31
GetGroups	32
Request Syntax	32
URI Request Parameters	32
Request Body	32
Response Syntax	32
Response Elements	33
Errors	33
See Also	34
GetIndexingRules	35
Request Syntax	35
URI Request Parameters	35
Request Body	35
Response Syntax	35
Response Elements	36
Errors	36
See Also	37
GetInsight	38
Request Syntax	38
URI Request Parameters	38
Request Body	38
Response Syntax	38
Response Elements	39
Errors	40
See Also	40
GetInsightEvents	41
Request Syntax	41
URI Request Parameters	41
Request Body	41
Response Syntax	42
Response Elements	43
Errors	43
See Also	44

GetInsightImpactGraph	45
Request Syntax	45
URI Request Parameters	45
Request Body	45
Response Syntax	46
Response Elements	47
Errors	48
See Also	48
GetInsightSummaries	50
Request Syntax	50
URI Request Parameters	50
Request Body	50
Response Syntax	52
Response Elements	53
Errors	53
See Also	54
GetRetrievedTracesGraph	55
Request Syntax	55
URI Request Parameters	55
Request Body	55
Response Syntax	56
Response Elements	58
Errors	59
See Also	59
GetSamplingRules	61
Request Syntax	61
URI Request Parameters	61
Request Body	61
Response Syntax	61
Response Elements	62
Errors	62
See Also	63
GetSamplingStatisticSummaries	64
Request Syntax	64
URI Request Parameters	64
Request Body	64

Response Syntax	64
Response Elements	65
Errors	65
See Also	65
GetSamplingTargets	67
Request Syntax	67
URI Request Parameters	67
Request Body	67
Response Syntax	68
Response Elements	68
Errors	69
See Also	69
GetServiceGraph	71
Request Syntax	71
URI Request Parameters	71
Request Body	71
Response Syntax	72
Response Elements	74
Errors	75
See Also	75
GetTimeSeriesServiceStatistics	77
Request Syntax	77
URI Request Parameters	77
Request Body	77
Response Syntax	79
Response Elements	80
Errors	81
See Also	81
GetTraceGraph	82
Request Syntax	82
URI Request Parameters	82
Request Body	82
Response Syntax	83
Response Elements	85
Errors	85
See Also	85

GetTraceSegmentDestination	87
Request Syntax	87
URI Request Parameters	87
Request Body	87
Response Syntax	87
Response Elements	87
Errors	88
See Also	88
GetTraceSummaries	90
Request Syntax	90
URI Request Parameters	90
Request Body	91
Response Syntax	92
Response Elements	96
Errors	96
See Also	97
ListResourcePolicies	98
Request Syntax	98
URI Request Parameters	98
Request Body	98
Response Syntax	98
Response Elements	99
Errors	99
See Also	100
ListRetrievedTraces	101
Request Syntax	101
URI Request Parameters	101
Request Body	101
Response Syntax	102
Response Elements	103
Errors	103
See Also	104
ListTagsForResource	105
Request Syntax	105
URI Request Parameters	105
Request Body	105

Response Syntax	106
Response Elements	106
Errors	106
See Also	107
PutEncryptionConfig	108
Request Syntax	108
URI Request Parameters	108
Request Body	108
Response Syntax	109
Response Elements	109
Errors	109
See Also	110
PutResourcePolicy	111
Request Syntax	111
URI Request Parameters	111
Request Body	111
Response Syntax	112
Response Elements	113
Errors	113
See Also	114
PutTelemetryRecords	115
Request Syntax	115
URI Request Parameters	115
Request Body	115
Response Syntax	116
Response Elements	116
Errors	116
See Also	117
PutTraceSegments	118
Request Syntax	119
URI Request Parameters	119
Request Body	119
Response Syntax	119
Response Elements	120
Errors	120
See Also	120

StartTraceRetrieval	122
Request Syntax	122
URI Request Parameters	122
Request Body	122
Response Syntax	123
Response Elements	123
Errors	124
See Also	124
TagResource	126
Request Syntax	126
URI Request Parameters	126
Request Body	126
Response Syntax	127
Response Elements	127
Errors	127
See Also	128
UntagResource	129
Request Syntax	129
URI Request Parameters	129
Request Body	129
Response Syntax	130
Response Elements	130
Errors	130
See Also	130
UpdateGroup	132
Request Syntax	132
URI Request Parameters	132
Request Body	132
Response Syntax	133
Response Elements	133
Errors	134
See Also	134
UpdateIndexingRule	136
Request Syntax	136
URI Request Parameters	136
Request Body	136

Response Syntax	137
Response Elements	137
Errors	137
See Also	138
UpdateSamplingRule	139
Request Syntax	139
URI Request Parameters	139
Request Body	139
Response Syntax	140
Response Elements	140
Errors	141
See Also	141
UpdateTraceSegmentDestination	142
Request Syntax	142
URI Request Parameters	142
Request Body	142
Response Syntax	142
Response Elements	143
Errors	143
See Also	143
Data Types	145
Alias	148
Contents	148
See Also	148
AnnotationValue	149
Contents	149
See Also	149
AnomalousService	150
Contents	150
See Also	150
AvailabilityZoneDetail	151
Contents	151
See Also	151
BackendConnectionErrors	152
Contents	152
See Also	153

Edge	154
Contents	154
See Also	155
EdgeStatistics	156
Contents	156
See Also	157
EncryptionConfig	158
Contents	158
See Also	158
ErrorRootCause	160
Contents	160
See Also	160
ErrorRootCauseEntity	161
Contents	161
See Also	161
ErrorRootCauseService	162
Contents	162
See Also	163
ErrorStatistics	164
Contents	164
See Also	164
FaultRootCause	165
Contents	165
See Also	165
FaultRootCauseEntity	166
Contents	166
See Also	166
FaultRootCauseService	167
Contents	167
See Also	168
FaultStatistics	169
Contents	169
See Also	169
ForecastStatistics	170
Contents	170
See Also	170

GraphLink	171
Contents	171
See Also	171
Group	172
Contents	172
See Also	173
GroupSummary	174
Contents	174
See Also	175
HistogramEntry	176
Contents	176
See Also	176
Http	177
Contents	177
See Also	178
IndexingRule	179
Contents	179
See Also	179
IndexingRuleValue	181
Contents	181
See Also	181
IndexingRuleValueUpdate	182
Contents	182
See Also	182
Insight	183
Contents	183
See Also	185
InsightEvent	186
Contents	186
See Also	187
InsightImpactGraphEdge	188
Contents	188
See Also	188
InsightImpactGraphService	189
Contents	189
See Also	190

InsightsConfiguration	191
Contents	191
See Also	191
InsightSummary	192
Contents	192
See Also	194
InstanceIdDetail	195
Contents	195
See Also	195
ProbabilisticRuleValue	196
Contents	196
See Also	196
ProbabilisticRuleValueUpdate	197
Contents	197
See Also	197
RequestImpactStatistics	198
Contents	198
See Also	198
ResourceARNDetail	199
Contents	199
See Also	199
ResourcePolicy	200
Contents	200
See Also	201
ResponseTimeRootCause	202
Contents	202
See Also	202
ResponseTimeRootCauseEntity	203
Contents	203
See Also	203
ResponseTimeRootCauseService	204
Contents	204
See Also	205
RetrievedService	206
Contents	206
See Also	206

RetrievedTrace	207
Contents	207
See Also	207
RootCauseException	209
Contents	209
See Also	209
SamplingRule	210
Contents	210
See Also	213
SamplingRuleRecord	214
Contents	214
See Also	214
SamplingRuleUpdate	215
Contents	215
See Also	217
SamplingStatisticsDocument	218
Contents	218
See Also	219
SamplingStatisticSummary	220
Contents	220
See Also	221
SamplingStrategy	222
Contents	222
See Also	222
SamplingTargetDocument	223
Contents	223
See Also	224
Segment	225
Contents	225
See Also	225
Service	226
Contents	226
See Also	228
ServiceId	229
Contents	229
See Also	229

ServiceStatistics	230
Contents	230
See Also	231
Span	232
Contents	232
See Also	232
Tag	233
Contents	233
See Also	234
TelemetryRecord	235
Contents	235
See Also	236
TimeSeriesServiceStatistics	237
Contents	237
See Also	238
Trace	239
Contents	239
See Also	240
TraceSummary	241
Contents	241
See Also	244
TraceUser	246
Contents	246
See Also	246
UnprocessedStatistics	247
Contents	247
See Also	247
UnprocessedTraceSegment	248
Contents	248
See Also	248
ValueWithServiceIds	249
Contents	249
See Also	249
Common Parameters	250
Common Errors	253

Welcome

AWS X-Ray provides APIs for managing debug traces and retrieving service maps and other data created by processing those traces.

This document was last published on August 31, 2025.

Actions

The following actions are supported:

- [BatchGetTraces](#)
- [CancelTraceRetrieval](#)
- [CreateGroup](#)
- [CreateSamplingRule](#)
- [DeleteGroup](#)
- [DeleteResourcePolicy](#)
- [DeleteSamplingRule](#)
- [GetEncryptionConfig](#)
- [GetGroup](#)
- [GetGroups](#)
- [GetIndexingRules](#)
- [GetInsight](#)
- [GetInsightEvents](#)
- [GetInsightImpactGraph](#)
- [GetInsightSummaries](#)
- [GetRetrievedTracesGraph](#)
- [GetSamplingRules](#)
- [GetSamplingStatisticSummaries](#)
- [GetSamplingTargets](#)
- [GetServiceGraph](#)
- [GetTimeSeriesServiceStatistics](#)
- [GetTraceGraph](#)
- [GetTraceSegmentDestination](#)
- [GetTraceSummaries](#)
- [ListResourcePolicies](#)
- [ListRetrievedTraces](#)
- [ListTagsForResource](#)

- [PutEncryptionConfig](#)
- [PutResourcePolicy](#)
- [PutTelemetryRecords](#)
- [PutTraceSegments](#)
- [StartTraceRetrieval](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateGroup](#)
- [UpdateIndexingRule](#)
- [UpdateSamplingRule](#)
- [UpdateTraceSegmentDestination](#)

BatchGetTraces

Note

You cannot find traces through this API if Transaction Search is enabled since trace is not indexed in X-Ray.

Retrieves a list of traces specified by ID. Each trace is a collection of segment documents that originates from a single request. Use `GetTraceSummaries` to get a list of trace IDs.

Request Syntax

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

{
  "NextToken": "string",
  "TraceIds": [ "string" ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

NextToken

Pagination token.

Type: String

Required: No

TraceIds

Specify the trace IDs of requests for which to retrieve segments.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 5 items.

Length Constraints: Minimum length of 1. Maximum length of 35.

Required: Yes

Response Syntax

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

{
  "NextToken": "string",
  "Traces": [
    {
      "Duration": number,
      "Id": "string",
      "LimitExceeded": boolean,
      "Segments": [
        {
          "Document": "string",
          "Id": "string"
        }
      ]
    }
  ],
  "UnprocessedTraceIds": [ "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

Pagination token.

Type: String

[Traces](#)

Full traces for the specified requests.

Type: Array of [Trace](#) objects

[UnprocessedTraceIds](#)

Trace IDs of requests that haven't been processed.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 35.

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

CancelTraceRetrieval

Cancels an ongoing trace retrieval job initiated by `StartTraceRetrieval` using the provided `RetrievalToken`. A successful cancellation will return an HTTP 200 response.

Request Syntax

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

{
  "RetrievalToken": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[RetrievalToken](#)

Retrieval token.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1020.

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](#).

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ResourceNotFoundException

The resource was not found. Verify that the name or Amazon Resource Name (ARN) of the resource is correct.

HTTP Status Code: 404

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

CreateGroup

Creates a group resource with a name and a filter expression.

Request Syntax

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

{
  "FilterExpression": "string",
  "GroupName": "string",
  "InsightsConfiguration": {
    "InsightsEnabled": boolean,
    "NotificationsEnabled": boolean
  },
  "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.

FilterExpression

The filter expression defining criteria by which to group traces.

Type: String

Required: No

GroupName

The case-sensitive name of the new group. Default is a reserved name and names must be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

Required: Yes

InsightsConfiguration

The structure containing configurations related to insights.

- The `InsightsEnabled` boolean can be set to true to enable insights for the new group or false to disable insights for the new group.
- The `NotificationsEnabled` boolean can be set to true to enable insights notifications for the new group. Notifications may only be enabled on a group with `InsightsEnabled` set to true.

Type: [InsightsConfiguration](#) object

Required: No

Tags

A map that contains one or more tag keys and tag values to attach to an X-Ray group. For more information about ways to use tags, see [Tagging AWS resources](#) in the *AWS General Reference*.

The following restrictions apply to tags:

- Maximum number of user-applied tags per resource: 50
- Maximum tag key length: 128 Unicode characters
- Maximum tag value length: 256 Unicode characters
- Valid values for key and value: a-z, A-Z, 0-9, space, and the following characters: `_` `:` `/` `=` `+` `-` and `@`
- Tag keys and values are case sensitive.
- Don't use `aws:` as a prefix for keys; it's reserved for AWS use.

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

{
  "Group": {
    "FilterExpression": "string",
    "GroupARN": "string",
    "GroupName": "string",
    "InsightsConfiguration": {
      "InsightsEnabled": boolean,
      "NotificationsEnabled": boolean
    }
  }
}
```

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.

Group

The group that was created. Contains the name of the group that was created, the Amazon Resource Name (ARN) of the group that was generated based on the group name, the filter expression, and the insight configuration that was assigned to the group.

Type: [Group](#) object

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

CreateSamplingRule

Creates a rule to control sampling behavior for instrumented applications. Services retrieve rules with [GetSamplingRules](#), and evaluate each rule in ascending order of *priority* for each request. If a rule matches, the service records a trace, borrowing it from the reservoir size. After 10 seconds, the service reports back to X-Ray with [GetSamplingTargets](#) to get updated versions of each in-use rule. The updated rule contains a trace quota that the service can use instead of borrowing from the reservoir.

Request Syntax

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

```
{
  "SamplingRule": {
    "Attributes": {
      "string" : "string"
    },
    "FixedRate": number,
    "Host": "string",
    "HTTPMethod": "string",
    "Priority": number,
    "ReservoirSize": number,
    "ResourceARN": "string",
    "RuleARN": "string",
    "RuleName": "string",
    "ServiceName": "string",
    "ServiceType": "string",
    "URLPath": "string",
    "Version": number
  },
  "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.

SamplingRule

The rule definition.

Type: [SamplingRule](#) object

Required: Yes

Tags

A map that contains one or more tag keys and tag values to attach to an X-Ray sampling rule. For more information about ways to use tags, see [Tagging AWS resources](#) in the *AWS General Reference*.

The following restrictions apply to tags:

- Maximum number of user-applied tags per resource: 50
- Maximum tag key length: 128 Unicode characters
- Maximum tag value length: 256 Unicode characters
- Valid values for key and value: a-z, A-Z, 0-9, space, and the following characters: _ . : / = + - and @
- Tag keys and values are case sensitive.
- Don't use `aws :` as a prefix for keys; it's reserved for AWS use.

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

```
{
  "SamplingRuleRecord": {
    "CreatedAt": number,
    "ModifiedAt": number,
    "SamplingRule": {
      "Attributes": {
        "string" : "string"
      },
      "FixedRate": number,
      "Host": "string",
      "HTTPMethod": "string",
      "Priority": number,
      "ReservoirSize": number,
      "ResourceARN": "string",
      "RuleARN": "string",
      "RuleName": "string",
      "ServiceName": "string",
      "ServiceType": "string",
      "URLPath": "string",
      "Version": 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.

SamplingRuleRecord

The saved rule definition and metadata.

Type: [SamplingRuleRecord](#) object

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

RuleLimitExceededException

You have reached the maximum number of sampling rules.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

DeleteGroup

Deletes a group resource.

Request Syntax

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

{
  "GroupARN": "string",
  "GroupName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

GroupARN

The ARN of the group that was generated on creation.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 400.

Required: No

GroupName

The case-sensitive name of the group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

Required: No

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](#).

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

DeleteResourcePolicy

Deletes a resource policy from the target AWS account.

Request Syntax

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

{
  "PolicyName": "string",
  "PolicyRevisionId": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

PolicyName

The name of the resource policy to delete.

Type: String

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

Pattern: `[\w+=, .@- ]+`

Required: Yes

PolicyRevisionId

Specifies a specific policy revision to delete. Provide a PolicyRevisionId to ensure an atomic delete operation. If the provided revision id does not match the latest policy revision id, an `InvalidPolicyRevisionIdException` exception is returned.

Type: String

Required: No

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](#).

InvalidPolicyRevisionIdException

A policy revision id was provided which does not match the latest policy revision. This exception is also if a policy revision id of 0 is provided via `PutResourcePolicy` and a policy with the same name already exists.

HTTP Status Code: 400

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

DeleteSamplingRule

Deletes a sampling rule.

Request Syntax

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

```
{
  "RuleARN": "string",
  "RuleName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

RuleARN

The ARN of the sampling rule. Specify a rule by either name or ARN, but not both.

Type: String

Required: No

RuleName

The name of the sampling rule. Specify a rule by either name or ARN, but not both.

Type: String

Required: No

Response Syntax

```
HTTP/1.1 200
```


Content-type: application/json

```
{
  "SamplingRuleRecord": {
    "CreatedAt": number,
    "ModifiedAt": number,
    "SamplingRule": {
      "Attributes": {
        "string" : "string"
      },
      "FixedRate": number,
      "Host": "string",
      "HTTPMethod": "string",
      "Priority": number,
      "ReservoirSize": number,
      "ResourceARN": "string",
      "RuleARN": "string",
      "RuleName": "string",
      "ServiceName": "string",
      "ServiceType": "string",
      "URLPath": "string",
      "Version": 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.

SamplingRuleRecord

The deleted rule definition and metadata.

Type: [SamplingRuleRecord](#) object

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetEncryptionConfig

Retrieves the current encryption configuration for X-Ray data.

Request Syntax

```
POST /EncryptionConfig 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
Content-type: application/json

{
  "EncryptionConfig": {
    "KeyId": "string",
    "Status": "string",
    "Type": "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.

[EncryptionConfig](#)

The encryption configuration document.

Type: [EncryptionConfig](#) object

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetGroup

Retrieves group resource details.

Request Syntax

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

{
  "GroupARN": "string",
  "GroupName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

GroupARN

The ARN of the group that was generated on creation.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 400.

Required: No

GroupName

The case-sensitive name of the group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

Required: No

Response Syntax

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

{
  "Group": {
    "FilterExpression": "string",
    "GroupARN": "string",
    "GroupName": "string",
    "InsightsConfiguration": {
      "InsightsEnabled": boolean,
      "NotificationsEnabled": boolean
    }
  }
}
```

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.

Group

The group that was requested. Contains the name of the group, the ARN of the group, the filter expression, and the insight configuration assigned to the group.

Type: [Group](#) object

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetGroups

Retrieves all active group details.

Request Syntax

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

{
  "NextToken": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

NextToken

Pagination token.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Required: No

Response Syntax

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

{
  "Groups": [
    {
      "FilterExpression": "string",
```



```
    "GroupARN": "string",
    "GroupName": "string",
    "InsightsConfiguration": {
        "InsightsEnabled": boolean,
        "NotificationsEnabled": boolean
    }
},
"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.

Groups

The collection of all active groups.

Type: Array of [GroupSummary](#) objects

NextToken

Pagination token.

Type: String

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetIndexingRules

Retrieves all indexing rules.

Indexing rules are used to determine the server-side sampling rate for spans ingested through the CloudWatchLogs destination and indexed by X-Ray. For more information, see [Transaction Search](#).

Request Syntax

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

{
  "NextToken": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[NextToken](#)

Specify the pagination token returned by a previous request to retrieve the next page of indexes.

Type: String

Required: No

Response Syntax

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

{
  "IndexingRules": [
```

```
{
  "ModifiedAt": number,
  "Name": "string",
  "Rule": { ... }
},
"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.

[IndexingRules](#)

Retrieves all indexing rules.

Type: Array of [IndexingRule](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10 items.

[NextToken](#)

Specify the pagination token returned by a previous request to retrieve the next page of indexes.

Type: String

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetInsight

Retrieves the summary information of an insight. This includes impact to clients and root cause services, the top anomalous services, the category, the state of the insight, and the start and end time of the insight.

Request Syntax

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

```
{
  "InsightId": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

InsightId

The insight's unique identifier. Use the GetInsightSummaries action to retrieve an InsightId.

Type: String

Pattern: `[0-9a-fA-F]{8}-[0-9a-fA-F]{4}-[1-5][0-9a-fA-F]{3}-[89abAB][0-9a-fA-F]{3}-[0-9a-fA-F]{12}`

Required: Yes

Response Syntax

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

```
{
```

```

"Insight": {
  "Categories": [ "string" ],
  "ClientRequestImpactStatistics": {
    "FaultCount": number,
    "OkCount": number,
    "TotalCount": number
  },
  "EndTime": number,
  "GroupARN": "string",
  "GroupName": "string",
  "InsightId": "string",
  "RootCauseServiceId": {
    "AccountId": "string",
    "Name": "string",
    "Names": [ "string" ],
    "Type": "string"
  },
  "RootCauseServiceRequestImpactStatistics": {
    "FaultCount": number,
    "OkCount": number,
    "TotalCount": number
  },
  "StartTime": number,
  "State": "string",
  "Summary": "string",
  "TopAnomalousServices": [
    {
      "ServiceId": {
        "AccountId": "string",
        "Name": "string",
        "Names": [ "string" ],
        "Type": "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.

[Insight](#)

The summary information of an insight.

Type: [Insight](#) object

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetInsightEvents

X-Ray reevaluates insights periodically until they're resolved, and records each intermediate state as an event. You can review an insight's events in the Impact Timeline on the Inspect page in the X-Ray console.

Request Syntax

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

{
  "InsightId": "string",
  "MaxResults": number,
  "NextToken": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

InsightId

The insight's unique identifier. Use the GetInsightSummaries action to retrieve an InsightId.

Type: String

Pattern: [0-9a-fA-F]{8}-[0-9a-fA-F]{4}-[1-5][0-9a-fA-F]{3}-[89abAB][0-9a-fA-F]{3}-[0-9a-fA-F]{12}

Required: Yes

MaxResults

Used to retrieve at most the specified value of events.

Type: Integer

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

Required: No

NextToken

Specify the pagination token returned by a previous request to retrieve the next page of events.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Required: No

Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "InsightEvents": [
    {
      "ClientRequestImpactStatistics": {
        "FaultCount": number,
        "OkCount": number,
        "TotalCount": number
      },
      "EventTime": number,
      "RootCauseServiceRequestImpactStatistics": {
        "FaultCount": number,
        "OkCount": number,
        "TotalCount": number
      },
      "Summary": "string",
      "TopAnomalousServices": [
        {
          "ServiceId": {
            "AccountId": "string",
            "Name": "string",
            "Names": [ "string" ],
            "Type": "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.

[InsightEvents](#)

A detailed description of the event. This includes the time of the event, client and root cause impact statistics, and the top anomalous service at the time of the event.

Type: Array of [InsightEvent](#) objects

[NextToken](#)

Use this token to retrieve the next page of insight events.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetInsightImpactGraph

Retrieves a service graph structure filtered by the specified insight. The service graph is limited to only structural information. For a complete service graph, use this API with the `GetServiceGraph` API.

Request Syntax

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

```
{
  "EndTime": number,
  "InsightId": "string",
  "NextToken": "string",
  "StartTime": number
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

EndTime

The estimated end time of the insight, in Unix time seconds. The `EndTime` is exclusive of the value provided. The time range between the start time and end time can't be more than six hours.

Type: Timestamp

Required: Yes

InsightId

The insight's unique identifier. Use the `GetInsightSummaries` action to retrieve an `InsightId`.

Type: String

Pattern: `[0-9a-fA-F]{8}-[0-9a-fA-F]{4}-[1-5][0-9a-fA-F]{3}-[89abAB][0-9a-fA-F]{3}-[0-9a-fA-F]{12}`

Required: Yes

[NextToken](#)

Specify the pagination token returned by a previous request to retrieve the next page of results.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Required: No

[StartTime](#)

The estimated start time of the insight, in Unix time seconds. The StartTime is inclusive of the value provided and can't be more than 30 days old.

Type: Timestamp

Required: Yes

Response Syntax

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

{
  "EndTime": number,
  "InsightId": "string",
  "NextToken": "string",
  "ServiceGraphEndTime": number,
  "ServiceGraphStartTime": number,
  "Services": [
    {
      "AccountId": "string",
      "Edges": [
        {
          "ReferenceId": number
        }
      ],
      "Name": "string",
```

```
    "Names": [ "string" ],
    "ReferenceId": number,
    "Type": "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 provided end time.

Type: Timestamp

InsightId

The insight's unique identifier.

Type: String

Pattern: [0-9a-fA-F]{8}-[0-9a-fA-F]{4}-[1-5][0-9a-fA-F]{3}-[89abAB][0-9a-fA-F]{3}-[0-9a-fA-F]{12}

NextToken

Pagination token.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

ServiceGraphEndTime

The time, in Unix seconds, at which the service graph ended.

Type: Timestamp

ServiceGraphStartTime

The time, in Unix seconds, at which the service graph started.

Type: Timestamp

Services

The AWS instrumented services related to the insight.

Type: Array of [InsightImpactGraphService](#) objects

StartTime

The provided start time.

Type: Timestamp

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetInsightSummaries

Retrieves the summaries of all insights in the specified group matching the provided filter values.

Request Syntax

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

```
{
  "EndTime": number,
  "GroupARN": "string",
  "GroupName": "string",
  "MaxResults": number,
  "NextToken": "string",
  "StartTime": number,
  "States": [ "string" ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

EndTime

The end of the time frame in which the insights ended. The end time can't be more than 30 days old.

Type: Timestamp

Required: Yes

GroupARN

The Amazon Resource Name (ARN) of the group. Required if the GroupName isn't provided.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 400.

Required: No

GroupName

The name of the group. Required if the GroupARN isn't provided.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

Required: No

MaxResults

The maximum number of results to display.

Type: Integer

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

Required: No

NextToken

Pagination token.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Required: No

StartTime

The beginning of the time frame in which the insights started. The start time can't be more than 30 days old.

Type: Timestamp

Required: Yes

States

The list of insight states.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 1 item.

Valid Values: ACTIVE | CLOSED

Required: No

Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "InsightSummaries": [
    {
      "Categories": [ "string" ],
      "ClientRequestImpactStatistics": {
        "FaultCount": number,
        "OkCount": number,
        "TotalCount": number
      },
      "EndTime": number,
      "GroupARN": "string",
      "GroupName": "string",
      "InsightId": "string",
      "LastUpdateTime": number,
      "RootCauseServiceId": {
        "AccountId": "string",
        "Name": "string",
        "Names": [ "string" ],
        "Type": "string"
      },
      "RootCauseServiceRequestImpactStatistics": {
        "FaultCount": number,
        "OkCount": number,
        "TotalCount": number
      },
      "StartTime": number,
      "State": "string",
      "Summary": "string",
      "TopAnomalousServices": [
        {
```

```
    "ServiceId": {
      "AccountId": "string",
      "Name": "string",
      "Names": [ "string" ],
      "Type": "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.

[InsightSummaries](#)

The summary of each insight within the group matching the provided filters. The summary contains the InsightID, start and end time, the root cause service, the root cause and client impact statistics, the top anomalous services, and the status of the insight.

Type: Array of [InsightSummary](#) objects

[NextToken](#)

Pagination token.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetRetrievedTracesGraph

Retrieves a service graph for traces based on the specified `RetrievalToken` from the CloudWatch log group generated by Transaction Search. This API does not initiate a retrieval job. You must first execute `StartTraceRetrieval` to obtain the required `RetrievalToken`.

The trace graph describes services that process incoming requests and any downstream services they call, which may include AWS resources, external APIs, or databases.

The response is empty until the `RetrievalStatus` is *COMPLETE*. Retry the request after the status changes from *RUNNING* or *SCHEDULED* to *COMPLETE* to access the full service graph.

When CloudWatch log is the destination, this API can support cross-account observability and service graph retrieval across linked accounts.

For retrieving graphs from X-Ray directly as opposed to the Transaction-Search Log group, see [GetTraceGraph](#).

Request Syntax

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

{
  "NextToken": "string",
  "RetrievalToken": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[NextToken](#)

Specify the pagination token returned by a previous request to retrieve the next page of indexes.

Type: String

Required: No

RetrievalToken

Retrieval token.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1020.

Required: Yes

Response Syntax

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

{
  "NextToken": "string",
  "RetrievalStatus": "string",
  "Services": [
    {
      "Links": [
        {
          "DestinationTraceIds": [ "string" ],
          "ReferenceType": "string",
          "SourceTraceId": "string"
        }
      ],
      "Service": {
        "AccountId": "string",
        "DurationHistogram": [
          {
            "Count": number,
            "Value": number
          }
        ],
        "Edges": [
          {
            "Aliases": [
              {
                "Name": "string",
                "Names": [ "string" ],
                "Type": "string"
              }
            ]
          }
        ]
      }
    }
  ]
}
```



```

    }
  ],
  "EdgeType": "string",
  "EndTime": number,
  "ReceivedEventAgeHistogram": [
    {
      "Count": number,
      "Value": number
    }
  ],
  "ReferenceId": number,
  "ResponseTimeHistogram": [
    {
      "Count": number,
      "Value": number
    }
  ],
  "StartTime": number,
  "SummaryStatistics": {
    "ErrorStatistics": {
      "OtherCount": number,
      "ThrottleCount": number,
      "TotalCount": number
    },
    "FaultStatistics": {
      "OtherCount": number,
      "TotalCount": number
    },
    "OkCount": number,
    "TotalCount": number,
    "TotalResponseTime": number
  }
}
],
"EndTime": number,
"Name": "string",
"Names": [ "string" ],
"ReferenceId": number,
"ResponseTimeHistogram": [
  {
    "Count": number,
    "Value": number
  }
]
],

```

```

    "Root": boolean,
    "StartTime": number,
    "State": "string",
    "SummaryStatistics": {
      "ErrorStatistics": {
        "OtherCount": number,
        "ThrottleCount": number,
        "TotalCount": number
      },
      "FaultStatistics": {
        "OtherCount": number,
        "TotalCount": number
      },
      "OkCount": number,
      "TotalCount": number,
      "TotalResponseTime": number
    },
    "Type": "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

Specify the pagination token returned by a previous request to retrieve the next page of indexes.

Type: String

RetrievalStatus

Status of the retrieval.

Type: String

Valid Values: SCHEDULED | RUNNING | COMPLETE | FAILED | CANCELLED | TIMEOUT

[Services](#)

Retrieved services.

Type: Array of [RetrievedService](#) objects

Array Members: Minimum number of 0 items. Maximum number of 1000 items.

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ResourceNotFoundException

The resource was not found. Verify that the name or Amazon Resource Name (ARN) of the resource is correct.

HTTP Status Code: 404

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetSamplingRules

Retrieves all sampling rules.

Request Syntax

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

```
{
  "NextToken": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

NextToken

Pagination token.

Type: String

Required: No

Response Syntax

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

```
{
  "NextToken": "string",
  "SamplingRuleRecords": [
    {
      "CreatedAt": number,
      "ModifiedAt": number,
```

```
"SamplingRule": {
  "Attributes": {
    "string" : "string"
  },
  "FixedRate": number,
  "Host": "string",
  "HTTPMethod": "string",
  "Priority": number,
  "ReservoirSize": number,
  "ResourceARN": "string",
  "RuleARN": "string",
  "RuleName": "string",
  "ServiceName": "string",
  "ServiceType": "string",
  "URLPath": "string",
  "Version": 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.

[NextToken](#)

Pagination token.

Type: String

[SamplingRuleRecords](#)

Rule definitions and metadata.

Type: Array of [SamplingRuleRecord](#) objects

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetSamplingStatisticSummaries

Retrieves information about recent sampling results for all sampling rules.

Request Syntax

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

```
{
  "NextToken": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

NextToken

Pagination token.

Type: String

Required: No

Response Syntax

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

{
  "NextToken": "string",
  "SamplingStatisticSummaries": [
    {
      "BorrowCount": number,
      "RequestCount": number,
      "RuleName": "string",
```



```
    "SampledCount": 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.

[NextToken](#)

Pagination token.

Type: String

[SamplingStatisticSummaries](#)

Information about the number of requests instrumented for each sampling rule.

Type: Array of [SamplingStatisticSummary](#) objects

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetSamplingTargets

Requests a sampling quota for rules that the service is using to sample requests.

Request Syntax

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

{
  "SamplingStatisticsDocuments": [
    {
      "BorrowCount": number,
      "ClientID": "string",
      "RequestCount": number,
      "RuleName": "string",
      "SampledCount": number,
      "Timestamp": number
    }
  ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

SamplingStatisticsDocuments

Information about rules that the service is using to sample requests.

Type: Array of [SamplingStatisticsDocument](#) objects

Array Members: Maximum number of 25 items.

Required: Yes

Response Syntax

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

{
  "LastRuleModification": number,
  "SamplingTargetDocuments": [
    {
      "FixedRate": number,
      "Interval": number,
      "ReservoirQuota": number,
      "ReservoirQuotaTTL": number,
      "RuleName": "string"
    }
  ],
  "UnprocessedStatistics": [
    {
      "ErrorCode": "string",
      "Message": "string",
      "RuleName": "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.

LastRuleModification

The last time a user changed the sampling rule configuration, in Unix time seconds. If the sampling rule configuration changed since the service last retrieved it, the service should call [GetSamplingRules](#) to get the latest version.

Type: Timestamp

SamplingTargetDocuments

Updated rules that the service should use to sample requests.

Type: Array of [SamplingTargetDocument](#) objects

UnprocessedStatistics

Information about [SamplingStatisticsDocument](#) that X-Ray could not process.

Type: Array of [UnprocessedStatistics](#) objects

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetServiceGraph

Retrieves a document that describes services that process incoming requests, and downstream services that they call as a result. Root services process incoming requests and make calls to downstream services. Root services are applications that use the [AWS X-Ray SDK](#). Downstream services can be other applications, AWS resources, HTTP web APIs, or SQL databases.

Request Syntax

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

{
  "EndTime": number,
  "GroupARN": "string",
  "GroupName": "string",
  "NextToken": "string",
  "StartTime": number
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

EndTime

The end of the timeframe for which to generate a graph, in Unix time seconds.

Type: Timestamp

Required: Yes

GroupARN

The Amazon Resource Name (ARN) of a group based on which you want to generate a graph.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 400.

Required: No

GroupName

The name of a group based on which you want to generate a graph.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

Required: No

NextToken

Pagination token.

Type: String

Required: No

StartTime

The start of the time frame for which to generate a graph, in Unix time seconds.

Type: Timestamp

Required: Yes

Response Syntax

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

{
  "ContainsOldGroupVersions": boolean,
  "EndTime": number,
  "NextToken": "string",
  "Services": [
    {
      "AccountId": "string",
      "DurationHistogram": [
        {
          "Count": number,
          "Value": number
        }
      ]
    }
  ],
```



```

    "Edges": [
      {
        "Aliases": [
          {
            "Name": "string",
            "Names": [ "string" ],
            "Type": "string"
          }
        ],
        "EdgeType": "string",
        "EndTime": number,
        "ReceivedEventAgeHistogram": [
          {
            "Count": number,
            "Value": number
          }
        ],
        "ReferenceId": number,
        "ResponseTimeHistogram": [
          {
            "Count": number,
            "Value": number
          }
        ],
        "StartTime": number,
        "SummaryStatistics": {
          "ErrorStatistics": {
            "OtherCount": number,
            "ThrottleCount": number,
            "TotalCount": number
          },
          "FaultStatistics": {
            "OtherCount": number,
            "TotalCount": number
          },
          "OkCount": number,
          "TotalCount": number,
          "TotalResponseTime": number
        }
      }
    ],
    "EndTime": number,
    "Name": "string",
    "Names": [ "string" ],

```

```

    "ReferenceId": number,
    "ResponseTimeHistogram": [
      {
        "Count": number,
        "Value": number
      }
    ],
    "Root": boolean,
    "StartTime": number,
    "State": "string",
    "SummaryStatistics": {
      "ErrorStatistics": {
        "OtherCount": number,
        "ThrottleCount": number,
        "TotalCount": number
      },
      "FaultStatistics": {
        "OtherCount": number,
        "TotalCount": number
      },
      "OkCount": number,
      "TotalCount": number,
      "TotalResponseTime": number
    },
    "Type": "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.

ContainsOldGroupVersions

A flag indicating whether the group's filter expression has been consistent, or if the returned service graph may show traces from an older version of the group's filter expression.

Type: Boolean

EndTime

The end of the time frame for which the graph was generated, in Unix time seconds.

Type: Timestamp

NextToken

Pagination token.

Type: String

Services

The services that have processed a traced request during the specified time frame.

Type: Array of [Service](#) objects

StartTime

The start of the time frame for which the graph was generated, in Unix time seconds.

Type: Timestamp

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetTimeSeriesServiceStatistics

Get an aggregation of service statistics defined by a specific time range.

Request Syntax

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

```
{
  "EndTime": number,
  "EntitySelectorExpression": "string",
  "ForecastStatistics": boolean,
  "GroupARN": "string",
  "GroupName": "string",
  "NextToken": "string",
  "Period": number,
  "StartTime": number
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

EndTime

The end of the time frame for which to aggregate statistics.

Type: Timestamp

Required: Yes

EntitySelectorExpression

A filter expression defining entities that will be aggregated for statistics. Supports ID, service, and edge functions. If no selector expression is specified, edge statistics are returned.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 500.

Required: No

ForecastStatistics

The forecasted high and low fault count values. Forecast enabled requests require the EntitySelectorExpression ID be provided.

Type: Boolean

Required: No

GroupARN

The Amazon Resource Name (ARN) of the group for which to pull statistics from.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 400.

Required: No

GroupName

The case-sensitive name of the group for which to pull statistics from.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

Required: No

NextToken

Pagination token.

Type: String

Required: No

Period

Aggregation period in seconds.

Type: Integer

Required: No

StartTime

The start of the time frame for which to aggregate statistics.

Type: Timestamp

Required: Yes

Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "ContainsOldGroupVersions": boolean,
  "NextToken": "string",
  "TimeSeriesServiceStatistics": [
    {
      "EdgeSummaryStatistics": {
        "ErrorStatistics": {
          "OtherCount": number,
          "ThrottleCount": number,
          "TotalCount": number
        },
        "FaultStatistics": {
          "OtherCount": number,
          "TotalCount": number
        },
        "OkCount": number,
        "TotalCount": number,
        "TotalResponseTime": number
      },
      "ResponseTimeHistogram": [
        {
          "Count": number,
          "Value": number
        }
      ],
      "ServiceForecastStatistics": {
        "FaultCountHigh": number,
        "FaultCountLow": number
      }
    }
  ]
}
```

```
    },
    "ServiceSummaryStatistics": {
      "ErrorStatistics": {
        "OtherCount": number,
        "ThrottleCount": number,
        "TotalCount": number
      },
      "FaultStatistics": {
        "OtherCount": number,
        "TotalCount": number
      },
      "OkCount": number,
      "TotalCount": number,
      "TotalResponseTime": 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.

ContainsOldGroupVersions

A flag indicating whether or not a group's filter expression has been consistent, or if a returned aggregation might show statistics from an older version of the group's filter expression.

Type: Boolean

NextToken

Pagination token.

Type: String

TimeSeriesServiceStatistics

The collection of statistics.

Type: Array of TimeSeriesServiceStatistics objects

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetTraceGraph

Retrieves a service graph for one or more specific trace IDs.

Request Syntax

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

{
  "NextToken": "string",
  "TraceIds": [ "string" ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

NextToken

Pagination token.

Type: String

Required: No

TraceIds

Trace IDs of requests for which to generate a service graph.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 5 items.

Length Constraints: Minimum length of 1. Maximum length of 35.

Required: Yes

Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "NextToken": "string",
  "Services": [
    {
      "AccountId": "string",
      "DurationHistogram": [
        {
          "Count": number,
          "Value": number
        }
      ],
      "Edges": [
        {
          "Aliases": [
            {
              "Name": "string",
              "Names": [ "string" ],
              "Type": "string"
            }
          ],
          "EdgeType": "string",
          "EndTime": number,
          "ReceivedEventAgeHistogram": [
            {
              "Count": number,
              "Value": number
            }
          ],
          "ReferenceId": number,
          "ResponseTimeHistogram": [
            {
              "Count": number,
              "Value": number
            }
          ],
          "StartTime": number,
          "SummaryStatistics": {
            "ErrorStatistics": {
```

```

        "OtherCount": number,
        "ThrottleCount": number,
        "TotalCount": number
    },
    "FaultStatistics": {
        "OtherCount": number,
        "TotalCount": number
    },
    "OkCount": number,
    "TotalCount": number,
    "TotalResponseTime": number
}
]
},
"EndTime": number,
"Name": "string",
"Names": [ "string" ],
"ReferenceId": number,
"ResponseTimeHistogram": [
    {
        "Count": number,
        "Value": number
    }
],
"Root": boolean,
"StartTime": number,
"State": "string",
"SummaryStatistics": {
    "ErrorStatistics": {
        "OtherCount": number,
        "ThrottleCount": number,
        "TotalCount": number
    },
    "FaultStatistics": {
        "OtherCount": number,
        "TotalCount": number
    },
    "OkCount": number,
    "TotalCount": number,
    "TotalResponseTime": number
},
"Type": "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](#)

Pagination token.

Type: String

[Services](#)

The services that have processed one of the specified requests.

Type: Array of [Service](#) objects

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetTraceSegmentDestination

Retrieves the current destination of data sent to PutTraceSegments and *OpenTelemetry protocol (OTLP)* endpoint. The Transaction Search feature requires a CloudWatchLogs destination. For more information, see [Transaction Search](#) and [OpenTelemetry](#).

Request Syntax

```
POST /GetTraceSegmentDestination 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
Content-type: application/json

{
  "Destination": "string",
  "Status": "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.

[Destination](#)

Retrieves the current destination.

Type: String

Valid Values: XRay | CloudWatchLogs

Status

Status of the retrieval.

Type: String

Valid Values: PENDING | ACTIVE

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

GetTraceSummaries

Retrieves IDs and annotations for traces available for a specified time frame using an optional filter. To get the full traces, pass the trace IDs to `BatchGetTraces`.

A filter expression can target traced requests that hit specific service nodes or edges, have errors, or come from a known user. For example, the following filter expression targets traces that pass through `api.example.com`:

```
service("api.example.com")
```

This filter expression finds traces that have an annotation named `account` with the value `12345`:

```
annotation.account = "12345"
```

For a full list of indexed fields and keywords that you can use in filter expressions, see [Use filter expressions](#) in the *AWS X-Ray Developer Guide*.

Request Syntax

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

```
{
  "EndTime": number,
  "FilterExpression": "string",
  "NextToken": "string",
  "Sampling": boolean,
  "SamplingStrategy": {
    "Name": "string",
    "Value": number
  },
  "StartTime": number,
  "TimeRangeType": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

EndTime

The end of the time frame for which to retrieve traces, in Unix time seconds.

Type: Timestamp

Required: Yes

FilterExpression

Specify a filter expression to retrieve trace summaries for services or requests that meet certain requirements.

Type: String

Required: No

NextToken

Specify the pagination token returned by a previous request to retrieve the next page of results.

Type: String

Required: No

Sampling

Set to `true` to get summaries for only a subset of available traces.

Type: Boolean

Required: No

SamplingStrategy

A parameter to indicate whether to enable sampling on trace summaries. Input parameters are Name and Value.

Type: [SamplingStrategy](#) object

Required: No

StartTime

The start of the time frame for which to retrieve traces, in Unix time seconds.

Type: Timestamp

Required: Yes

TimeRangeType

Query trace summaries by TraceId (trace start time), Event (trace update time), or Service (trace segment end time).

Type: String

Valid Values: TraceId | Event | Service

Required: No

Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "ApproximateTime": number,
  "NextToken": "string",
  "TracesProcessedCount": number,
  "TraceSummaries": [
    {
      "Annotations": {
        "string" : [
          {
            "AnnotationValue": {
              "BooleanValue": boolean,
              "NumberValue": number,
              "StringValue": "string"
            },
            "ServiceIds": [
              {
                "AccountId": "string",
                "Name": "string",
                "Names": [ "string" ],
                "Type": "string"
              }
            ]
          }
        ]
      }
    }
  ]
}
```

```

    }
  ]
}
],
"AvailabilityZones": [
  {
    "Name": "string"
  }
],
"Duration": number,
"EntryPoint": {
  "AccountId": "string",
  "Name": "string",
  "Names": [ "string" ],
  "Type": "string"
},
"ErrorRootCauses": [
  {
    "ClientImpacting": boolean,
    "Services": [
      {
        "AccountId": "string",
        "EntityPath": [
          {
            "Exceptions": [
              {
                "Message": "string",
                "Name": "string"
              }
            ],
            "Name": "string",
            "Remote": boolean
          }
        ],
        "Inferred": boolean,
        "Name": "string",
        "Names": [ "string" ],
        "Type": "string"
      }
    ]
  }
],
"FaultRootCauses": [

```

```

    {
      "ClientImpacting": boolean,
      "Services": [
        {
          "AccountId": "string",
          "EntityPath": [
            {
              "Exceptions": [
                {
                  "Message": "string",
                  "Name": "string"
                }
              ],
              "Name": "string",
              "Remote": boolean
            }
          ],
          "Inferred": boolean,
          "Name": "string",
          "Names": [ "string " ],
          "Type": "string"
        }
      ]
    },
    "HasError": boolean,
    "HasFault": boolean,
    "HasThrottle": boolean,
    "Http": {
      "ClientIp": "string",
      "HttpMethod": "string",
      "HttpStatus": number,
      "HttpURL": "string",
      "UserAgent": "string"
    },
    "Id": "string",
    "InstanceIds": [
      {
        "Id": "string"
      }
    ],
    "IsPartial": boolean,
    "MatchedEventTime": number,
    "ResourceARNs": [

```

```

    {
      "ARN": "string"
    }
  ],
  "ResponseTime": number,
  "ResponseTimeRootCauses": [
    {
      "ClientImpacting": boolean,
      "Services": [
        {
          "AccountId": "string",
          "EntityPath": [
            {
              "Coverage": number,
              "Name": "string",
              "Remote": boolean
            }
          ],
          "Inferred": boolean,
          "Name": "string",
          "Names": [ "string" ],
          "Type": "string"
        }
      ]
    }
  ],
  "Revision": number,
  "ServiceIds": [
    {
      "AccountId": "string",
      "Name": "string",
      "Names": [ "string" ],
      "Type": "string"
    }
  ],
  "StartTime": number,
  "Users": [
    {
      "ServiceIds": [
        {
          "AccountId": "string",
          "Name": "string",
          "Names": [ "string" ],
          "Type": "string"
        }
      ]
    }
  ]
}

```

```
    }  
  ],  
  "UserName": "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.

ApproximateTime

The start time of this page of results.

Type: Timestamp

NextToken

If the requested time frame contained more than one page of results, you can use this token to retrieve the next page. The first page contains the most recent results, closest to the end of the time frame.

Type: String

TracesProcessedCount

The total number of traces processed, including traces that did not match the specified filter expression.

Type: Long

TraceSummaries

Trace IDs and annotations for traces that were found in the specified time frame.

Type: Array of [TraceSummary](#) objects

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

ListResourcePolicies

Returns the list of resource policies in the target AWS account.

Request Syntax

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

```
{
  "NextToken": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

NextToken

Not currently supported.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Required: No

Response Syntax

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

```
{
  "NextToken": "string",
  "ResourcePolicies": [
    {
```

```
    "LastUpdatedTime": number,  
    "PolicyDocument": "string",  
    "PolicyName": "string",  
    "PolicyRevisionId": "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](#)

Pagination token. Not currently supported.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

[ResourcePolicies](#)

The list of resource policies in the target AWS account.

Type: Array of [ResourcePolicy](#) objects

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

ListRetrievedTraces

Retrieves a list of traces for a given `RetrievalToken` from the CloudWatch log group generated by Transaction Search. For information on what each trace returns, see [BatchGetTraces](#).

This API does not initiate a retrieval process. To start a trace retrieval, use `StartTraceRetrieval`, which generates the required `RetrievalToken`.

When the `RetrievalStatus` is not *COMPLETE*, the API will return an empty response. Retry the request once the retrieval has completed to access the full list of traces.

For cross-account observability, this API can retrieve traces from linked accounts when CloudWatch log is set as the destination across relevant accounts. For more details, see [CloudWatch cross-account observability](#).

For retrieving data from X-Ray directly as opposed to the Transaction Search generated log group, see [BatchGetTraces](#).

Request Syntax

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

{
  "NextToken": "string",
  "RetrievalToken": "string",
  "TraceFormat": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[NextToken](#)

Specify the pagination token returned by a previous request to retrieve the next page of indexes.

Type: String

Required: No

RetrievalToken

Retrieval token.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1020.

Required: Yes

TraceFormat

Format of the requested traces.

Type: String

Valid Values: XRAY | OTEL

Required: No

Response Syntax

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

{
  "NextToken": "string",
  "RetrievalStatus": "string",
  "TraceFormat": "string",
  "Traces": [
    {
      "Duration": number,
      "Id": "string",
      "Spans": [
        {
          "Document": "string",
          "Id": "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](#)

Specify the pagination token returned by a previous request to retrieve the next page of indexes.

Type: String

[RetrievalStatus](#)

Status of the retrieval.

Type: String

Valid Values: SCHEDULED | RUNNING | COMPLETE | FAILED | CANCELLED | TIMEOUT

[TraceFormat](#)

Format of the requested traces.

Type: String

Valid Values: XRAY | OTEL

[Traces](#)

Full traces for the specified requests.

Type: Array of [RetrievedTrace](#) objects

Array Members: Minimum number of 0 items. Maximum number of 5 items.

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ResourceNotFoundException

The resource was not found. Verify that the name or Amazon Resource Name (ARN) of the resource is correct.

HTTP Status Code: 404

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

ListTagsForResource

Returns a list of tags that are applied to the specified AWS X-Ray group or sampling rule.

Request Syntax

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

```
{
  "NextToken": "string",
  "ResourceARN": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

NextToken

A pagination token. If multiple pages of results are returned, use the `NextToken` value returned with the current page of results as the value of this parameter to get the next page of results.

Type: String

Required: No

ResourceARN

The Amazon Resource Number (ARN) of an X-Ray group or sampling rule.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1011.

Required: Yes

Response Syntax

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

{
  "NextToken": "string",
  "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.

[NextToken](#)

A pagination token. If multiple pages of results are returned, use the `NextToken` value returned with the current page of results to get the next page of results.

Type: String

[Tags](#)

A list of tags, as key and value pairs, that is associated with the specified X-Ray group or sampling rule.

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](#).

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ResourceNotFoundException

The resource was not found. Verify that the name or Amazon Resource Name (ARN) of the resource is correct.

HTTP Status Code: 404

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

PutEncryptionConfig

Updates the encryption configuration for X-Ray data.

Request Syntax

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

{
  "KeyId": "string",
  "Type": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

KeyId

An AWS KMS key in one of the following formats:

- **Alias** - The name of the key. For example, `alias/MyKey`.
- **Key ID** - The KMS key ID of the key. For example, `ae4aa6d49-a4d8-9df9-a475-4ff6d7898456`. AWS X-Ray does not support asymmetric KMS keys.
- **ARN** - The full Amazon Resource Name of the key ID or alias. For example, `arn:aws:kms:us-east-2:123456789012:key/ae4aa6d49-a4d8-9df9-a475-4ff6d7898456`. Use this format to specify a key in a different account.

Omit this key if you set Type to NONE.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 3000.

Required: No

Type

The type of encryption. Set to KMS to use your own key for encryption. Set to NONE for default encryption.

Type: String

Valid Values: NONE | KMS

Required: Yes

Response Syntax

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

{
  "EncryptionConfig": {
    "KeyId": "string",
    "Status": "string",
    "Type": "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.

[EncryptionConfig](#)

The new encryption configuration.

Type: [EncryptionConfig](#) object

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

PutResourcePolicy

Sets the resource policy to grant one or more AWS services and accounts permissions to access X-Ray. Each resource policy will be associated with a specific AWS account. Each AWS account can have a maximum of 5 resource policies, and each policy name must be unique within that account. The maximum size of each resource policy is 5KB.

Request Syntax

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

{
  "BypassPolicyLockoutCheck": boolean,
  "PolicyDocument": "string",
  "PolicyName": "string",
  "PolicyRevisionId": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[BypassPolicyLockoutCheck](#)

A flag to indicate whether to bypass the resource policy lockout safety check.

Important

Setting this value to true increases the risk that the policy becomes unmanageable. Do not set this value to true indiscriminately.

Use this parameter only when you include a policy in the request and you intend to prevent the principal that is making the request from making a subsequent PutResourcePolicy request.

The default value is false.

Type: Boolean

Required: No

PolicyDocument

The resource policy document, which can be up to 5kb in size.

Type: String

Required: Yes

PolicyName

The name of the resource policy. Must be unique within a specific AWS account.

Type: String

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

Pattern: `[\w+=, .@- ]+`

Required: Yes

PolicyRevisionId

Specifies a specific policy revision, to ensure an atomic create operation. By default the resource policy is created if it does not exist, or updated with an incremented revision id. The revision id is unique to each policy in the account.

If the policy revision id does not match the latest revision id, the operation will fail with an `InvalidPolicyRevisionIdException` exception. You can also provide a `PolicyRevisionId` of 0. In this case, the operation will fail with an `InvalidPolicyRevisionIdException` exception if a resource policy with the same name already exists.

Type: String

Required: No

Response Syntax

```
HTTP/1.1 200
```



```
Content-type: application/json

{
  "ResourcePolicy": {
    "LastUpdatedTime": number,
    "PolicyDocument": "string",
    "PolicyName": "string",
    "PolicyRevisionId": "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.

[ResourcePolicy](#)

The resource policy document, as provided in the PutResourcePolicyRequest.

Type: [ResourcePolicy](#) object

Errors

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

InvalidPolicyRevisionIdException

A policy revision id was provided which does not match the latest policy revision. This exception is also if a policy revision id of 0 is provided via PutResourcePolicy and a policy with the same name already exists.

HTTP Status Code: 400

LockoutPreventionException

The provided resource policy would prevent the caller of this request from calling PutResourcePolicy in the future.

HTTP Status Code: 400

MalformedPolicyDocumentException

Invalid policy document provided in request.

HTTP Status Code: 400

PolicyCountLimitExceededException

Exceeded the maximum number of resource policies for a target AWS account.

HTTP Status Code: 400

PolicySizeLimitExceededException

Exceeded the maximum size for a resource policy.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

PutTelemetryRecords

Used by the AWS X-Ray daemon to upload telemetry.

Request Syntax

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

{
  "EC2InstanceId": "string",
  "Hostname": "string",
  "ResourceARN": "string",
  "TelemetryRecords": [
    {
      "BackendConnectionErrors": {
        "ConnectionRefusedCount": number,
        "HTTPCode4XXCount": number,
        "HTTPCode5XXCount": number,
        "OtherCount": number,
        "TimeoutCount": number,
        "UnknownHostCount": number
      },
      "SegmentsReceivedCount": number,
      "SegmentsRejectedCount": number,
      "SegmentsSentCount": number,
      "SegmentsSpilloverCount": number,
      "Timestamp": number
    }
  ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

EC2InstanceId

Type: String

Length Constraints: Maximum length of 20.

Required: No

[Hostname](#)

Type: String

Length Constraints: Maximum length of 255.

Required: No

[ResourceARN](#)

Type: String

Length Constraints: Maximum length of 500.

Required: No

[TelemetryRecords](#)

Type: Array of [TelemetryRecord](#) objects

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](#).

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

PutTraceSegments

Uploads segment documents to AWS X-Ray. A segment document can be a completed segment, an in-progress segment, or an array of subsegments.

Segments must include the following fields. For the full segment document schema, see [AWS X-Ray Segment Documents](#) in the *AWS X-Ray Developer Guide*.

Required segment document fields

- `name` - The name of the service that handled the request.
- `id` - A 64-bit identifier for the segment, unique among segments in the same trace, in 16 hexadecimal digits.
- `trace_id` - A unique identifier that connects all segments and subsegments originating from a single client request.
- `start_time` - Time the segment or subsegment was created, in floating point seconds in epoch time, accurate to milliseconds. For example, 1480615200.010 or 1.480615200010E9.
- `end_time` - Time the segment or subsegment was closed. For example, 1480615200.090 or 1.480615200090E9. Specify either an `end_time` or `in_progress`.
- `in_progress` - Set to `true` instead of specifying an `end_time` to record that a segment has been started, but is not complete. Send an in-progress segment when your application receives a request that will take a long time to serve, to trace that the request was received. When the response is sent, send the complete segment to overwrite the in-progress segment.

A `trace_id` consists of three numbers separated by hyphens. For example, 1-58406520-a006649127e371903a2de979. For trace IDs created by an X-Ray SDK, or by AWS services integrated with X-Ray, a trace ID includes:

Trace ID Format

- The version number, for instance, 1.
- The time of the original request, in Unix epoch time, in 8 hexadecimal digits. For example, 10:00AM December 2nd, 2016 PST in epoch time is 1480615200 seconds, or 58406520 in hexadecimal.
- A 96-bit identifier for the trace, globally unique, in 24 hexadecimal digits.

Note

Trace IDs created via OpenTelemetry have a different format based on the [W3C Trace Context specification](#). A W3C trace ID must be formatted in the X-Ray trace ID format when sending to X-Ray. For example, a W3C trace ID 4efaaf4d1e8720b39541901950019ee5 should be formatted as 1-4efaaf4d-1e8720b39541901950019ee5 when sending to X-Ray. While X-Ray trace IDs include the original request timestamp in Unix epoch time, this is not required or validated.

Request Syntax

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

{
  "TraceSegmentDocuments": [ "string" ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[TraceSegmentDocuments](#)

A string containing a JSON document defining one or more segments or subsegments.

Type: Array of strings

Required: Yes

Response Syntax

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

```
{
  "UnprocessedTraceSegments": [
    {
      "ErrorCode": "string",
      "Id": "string",
      "Message": "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.

UnprocessedTraceSegments

Segments that failed processing.

Type: Array of [UnprocessedTraceSegment](#) objects

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

StartTraceRetrieval

Initiates a trace retrieval process using the specified time range and for the given trace IDs in the Transaction Search generated CloudWatch log group. For more information, see [Transaction Search](#).

API returns a `RetrievalToken`, which can be used with `ListRetrievedTraces` or `GetRetrievedTracesGraph` to fetch results. Retrievals will time out after 60 minutes. To execute long time ranges, consider segmenting into multiple retrievals.

If you are using [CloudWatch cross-account observability](#), you can use this operation in a monitoring account to retrieve data from a linked source account, as long as both accounts have transaction search enabled.

For retrieving data from X-Ray directly as opposed to the Transaction-Search Log group, see [BatchGetTraces](#).

Request Syntax

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

{
  "EndTime": number,
  "StartTime": number,
  "TraceIds": [ "string" ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[EndTime](#)

The end of the time range to retrieve traces. The range is inclusive, so the specified end time is included in the query. Specified as epoch time, the number of seconds since January 1, 1970, 00:00:00 UTC.

Type: Timestamp

Required: Yes

StartTime

The start of the time range to retrieve traces. The range is inclusive, so the specified start time is included in the query. Specified as epoch time, the number of seconds since January 1, 1970, 00:00:00 UTC.

Type: Timestamp

Required: Yes

TraceIds

Specify the trace IDs of the traces to be retrieved.

Type: Array of strings

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

Length Constraints: Minimum length of 1. Maximum length of 35.

Required: Yes

Response Syntax

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

{
  "RetrievalToken": "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.

RetrievalToken

Retrieval token.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1020.

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ResourceNotFoundException

The resource was not found. Verify that the name or Amazon Resource Name (ARN) of the resource is correct.

HTTP Status Code: 404

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

TagResource

Applies tags to an existing AWS X-Ray group or sampling rule.

Request Syntax

```
POST /TagResource 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 Number (ARN) of an X-Ray group or sampling rule.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1011.

Required: Yes

Tags

A map that contains one or more tag keys and tag values to attach to an X-Ray group or sampling rule. For more information about ways to use tags, see [Tagging AWS resources](#) in the *AWS General Reference*.

The following restrictions apply to tags:

- Maximum number of user-applied tags per resource: 50
- Maximum tag key length: 128 Unicode characters
- Maximum tag value length: 256 Unicode characters
- Valid values for key and value: a-z, A-Z, 0-9, space, and the following characters: _ . : / = + - and @
- Tag keys and values are case sensitive.
- Don't use `aws :` as a prefix for keys; it's reserved for AWS use. You cannot edit or delete system tags.

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](#).

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ResourceNotFoundException

The resource was not found. Verify that the name or Amazon Resource Name (ARN) of the resource is correct.

HTTP Status Code: 404

ThrottledException

The request exceeds the maximum number of requests per second.

HTTP Status Code: 429

TooManyTagsException

You have exceeded the maximum number of tags you can apply to this resource.

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](#)

UntagResource

Removes tags from an AWS X-Ray group or sampling rule. You cannot edit or delete system tags (those with an `aws:` prefix).

Request Syntax

```
POST /UntagResource 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 Number (ARN) of an X-Ray group or sampling rule.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1011.

Required: Yes

TagKeys

Keys for one or more tags that you want to remove from an X-Ray group or sampling rule.

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](#).

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ResourceNotFoundException

The resource was not found. Verify that the name or Amazon Resource Name (ARN) of the resource is correct.

HTTP Status Code: 404

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

UpdateGroup

Updates a group resource.

Request Syntax

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

{
  "FilterExpression": "string",
  "GroupARN": "string",
  "GroupName": "string",
  "InsightsConfiguration": {
    "InsightsEnabled": boolean,
    "NotificationsEnabled": boolean
  }
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

FilterExpression

The updated filter expression defining criteria by which to group traces.

Type: String

Required: No

GroupARN

The ARN that was generated upon creation.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 400.

Required: No

GroupName

The case-sensitive name of the group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

Required: No

InsightsConfiguration

The structure containing configurations related to insights.

- The `InsightsEnabled` boolean can be set to true to enable insights for the group or false to disable insights for the group.
- The `NotificationsEnabled` boolean can be set to true to enable insights notifications for the group. Notifications can only be enabled on a group with `InsightsEnabled` set to true.

Type: [InsightsConfiguration](#) object

Required: No

Response Syntax

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

{
  "Group": {
    "FilterExpression": "string",
    "GroupARN": "string",
    "GroupName": "string",
    "InsightsConfiguration": {
      "InsightsEnabled": boolean,
      "NotificationsEnabled": boolean
    }
  }
}
```

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.

Group

The group that was updated. Contains the name of the group that was updated, the ARN of the group that was updated, the updated filter expression, and the updated insight configuration assigned to the group.

Type: [Group](#) object

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

UpdateIndexingRule

Modifies an indexing rule's configuration.

Indexing rules are used for determining the sampling rate for spans indexed from CloudWatch Logs. For more information, see [Transaction Search](#).

Request Syntax

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

{
  "Name": "string",
  "Rule": { ... }
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

Name

Name of the indexing rule to be updated.

Type: String

Required: Yes

Rule

Rule configuration to be updated.

Type: [IndexingRuleValueUpdate](#) object

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

Required: Yes

Response Syntax

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

{
  "IndexingRule": {
    "ModifiedAt": number,
    "Name": "string",
    "Rule": { ... }
  }
}
```

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.

[IndexingRule](#)

Updated indexing rule.

Type: [IndexingRule](#) object

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ResourceNotFoundException

The resource was not found. Verify that the name or Amazon Resource Name (ARN) of the resource is correct.

HTTP Status Code: 404

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

UpdateSamplingRule

Modifies a sampling rule's configuration.

Request Syntax

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

```
{
  "SamplingRuleUpdate": {
    "Attributes": {
      "string" : "string"
    },
    "FixedRate": number,
    "Host": "string",
    "HTTPMethod": "string",
    "Priority": number,
    "ReservoirSize": number,
    "ResourceARN": "string",
    "RuleARN": "string",
    "RuleName": "string",
    "ServiceName": "string",
    "ServiceType": "string",
    "URLPath": "string"
  }
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

SamplingRuleUpdate

The rule and fields to change.

Type: [SamplingRuleUpdate](#) object

Required: Yes

Response Syntax

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

{
  "SamplingRuleRecord": {
    "CreatedAt": number,
    "ModifiedAt": number,
    "SamplingRule": {
      "Attributes": {
        "string" : "string"
      },
      "FixedRate": number,
      "Host": "string",
      "HTTPMethod": "string",
      "Priority": number,
      "ReservoirSize": number,
      "ResourceARN": "string",
      "RuleARN": "string",
      "RuleName": "string",
      "ServiceName": "string",
      "ServiceType": "string",
      "URLPath": "string",
      "Version": 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.

SamplingRuleRecord

The updated rule definition and metadata.

Type: [SamplingRuleRecord](#) object

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

UpdateTraceSegmentDestination

Modifies the destination of data sent to PutTraceSegments. The Transaction Search feature requires the CloudWatchLogs destination. For more information, see [Transaction Search](#).

Request Syntax

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

{
  "Destination": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[Destination](#)

The configured destination of trace segments.

Type: String

Valid Values: XRay | CloudWatchLogs

Required: No

Response Syntax

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

{
  "Destination": "string",
  "Status": "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.

Destination

The destination of the trace segments.

Type: String

Valid Values: XRay | CloudWatchLogs

Status

The status of the update.

Type: String

Valid Values: PENDING | ACTIVE

Errors

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

InvalidRequestException

The request is missing required parameters or has invalid parameters.

HTTP Status Code: 400

ThrottledException

The request exceeds the maximum number of requests per second.

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](#)

Data Types

The AWS X-Ray 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:

- [Alias](#)
- [AnnotationValue](#)
- [AnomalousService](#)
- [AvailabilityZoneDetail](#)
- [BackendConnectionErrors](#)
- [Edge](#)
- [EdgeStatistics](#)
- [EncryptionConfig](#)
- [ErrorRootCause](#)
- [ErrorRootCauseEntity](#)
- [ErrorRootCauseService](#)
- [ErrorStatistics](#)
- [FaultRootCause](#)
- [FaultRootCauseEntity](#)
- [FaultRootCauseService](#)
- [FaultStatistics](#)
- [ForecastStatistics](#)
- [GraphLink](#)
- [Group](#)
- [GroupSummary](#)

- [HistogramEntry](#)
- [Http](#)
- [IndexingRule](#)
- [IndexingRuleValue](#)
- [IndexingRuleValueUpdate](#)
- [Insight](#)
- [InsightEvent](#)
- [InsightImpactGraphEdge](#)
- [InsightImpactGraphService](#)
- [InsightsConfiguration](#)
- [InsightSummary](#)
- [InstanceIdDetail](#)
- [ProbabilisticRuleValue](#)
- [ProbabilisticRuleValueUpdate](#)
- [RequestImpactStatistics](#)
- [ResourceARNDetail](#)
- [ResourcePolicy](#)
- [ResponseTimeRootCause](#)
- [ResponseTimeRootCauseEntity](#)
- [ResponseTimeRootCauseService](#)
- [RetrievedService](#)
- [RetrievedTrace](#)
- [RootCauseException](#)
- [SamplingRule](#)
- [SamplingRuleRecord](#)
- [SamplingRuleUpdate](#)
- [SamplingStatisticsDocument](#)
- [SamplingStatisticSummary](#)
- [SamplingStrategy](#)
- [SamplingTargetDocument](#)

- [Segment](#)
- [Service](#)
- [ServiceId](#)
- [ServiceStatistics](#)
- [Span](#)
- [Tag](#)
- [TelemetryRecord](#)
- [TimeSeriesServiceStatistics](#)
- [Trace](#)
- [TraceSummary](#)
- [TraceUser](#)
- [UnprocessedStatistics](#)
- [UnprocessedTraceSegment](#)
- [ValueWithServiceIds](#)

Alias

An alias for an edge.

Contents

Name

The canonical name of the alias.

Type: String

Required: No

Names

A list of names for the alias, including the canonical name.

Type: Array of strings

Required: No

Type

The type of the alias.

Type: String

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](#)

AnnotationValue

Value of a segment annotation. Has one of three value types: Number, Boolean, or String.

Contents

BooleanValue

Value for a Boolean annotation.

Type: Boolean

Required: No

NumberValue

Value for a Number annotation.

Type: Double

Required: No

StringValue

Value for a String annotation.

Type: String

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](#)

AnomalousService

The service within the service graph that has anomalously high fault rates.

Contents

ServiceId

Type: [ServiceId](#) 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](#)

AvailabilityZoneDetail

A list of Availability Zones corresponding to the segments in a trace.

Contents

Name

The name of a corresponding Availability Zone.

Type: String

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](#)

BackendConnectionErrors

Contents

ConnectionRefusedCount

Type: Integer

Required: No

HTTPCode4XXCount

Type: Integer

Required: No

HTTPCode5XXCount

Type: Integer

Required: No

OtherCount

Type: Integer

Required: No

TimeoutCount

Type: Integer

Required: No

UnknownHostCount

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](#)

Edge

Information about a connection between two services. An edge can be a synchronous connection, such as typical call between client and service, or an asynchronous link, such as a Lambda function which retrieves an event from an SNS queue.

Contents

Aliases

Aliases for the edge.

Type: Array of [Alias](#) objects

Required: No

EdgeType

Describes an asynchronous connection, with a value of `link`.

Type: String

Required: No

EndTime

The end time of the last segment on the edge.

Type: Timestamp

Required: No

ReceivedEventAgeHistogram

A histogram that maps the spread of event age when received by consumers. Age is calculated each time an event is received. Only populated when *EdgeType* is `link`.

Type: Array of [HistogramEntry](#) objects

Required: No

ReferenceId

Identifier of the edge. Unique within a service map.

Type: Integer

Required: No

ResponseTimeHistogram

A histogram that maps the spread of client response times on an edge. Only populated for synchronous edges.

Type: Array of [HistogramEntry](#) objects

Required: No

StartTime

The start time of the first segment on the edge.

Type: Timestamp

Required: No

SummaryStatistics

Response statistics for segments on the edge.

Type: [EdgeStatistics](#) 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](#)

EdgeStatistics

Response statistics for an edge.

Contents

ErrorStatistics

Information about requests that failed with a 4xx Client Error status code.

Type: [ErrorStatistics](#) object

Required: No

FaultStatistics

Information about requests that failed with a 5xx Server Error status code.

Type: [FaultStatistics](#) object

Required: No

OkCount

The number of requests that completed with a 2xx Success status code.

Type: Long

Required: No

TotalCount

The total number of completed requests.

Type: Long

Required: No

TotalResponseTime

The aggregate response time of completed requests.

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](#)

EncryptionConfig

A configuration document that specifies encryption configuration settings.

Contents

KeyId

The ID of the KMS key used for encryption, if applicable.

Type: String

Required: No

Status

The encryption status. While the status is UPDATING, X-Ray may encrypt data with a combination of the new and old settings.

Type: String

Valid Values: UPDATING | ACTIVE

Required: No

Type

The type of encryption. Set to KMS for encryption with KMS keys. Set to NONE for default encryption.

Type: String

Valid Values: NONE | KMS

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](#)

ErrorRootCause

The root cause of a trace summary error.

Contents

ClientImpacting

A flag that denotes that the root cause impacts the trace client.

Type: Boolean

Required: No

Services

A list of services corresponding to an error. A service identifies a segment and it contains a name, account ID, type, and inferred flag.

Type: Array of [ErrorRootCauseService](#) 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](#)

ErrorRootCauseEntity

A collection of segments and corresponding subsegments associated to a trace summary error.

Contents

Exceptions

The types and messages of the exceptions.

Type: Array of [RootCauseException](#) objects

Required: No

Name

The name of the entity.

Type: String

Required: No

Remote

A flag that denotes a remote subsegment.

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](#)

ErrorRootCauseService

A collection of fields identifying the services in a trace summary error.

Contents

AccountId

The account ID associated to the service.

Type: String

Required: No

EntityPath

The path of root cause entities found on the service.

Type: Array of [ErrorRootCauseEntity](#) objects

Required: No

Inferred

A Boolean value indicating if the service is inferred from the trace.

Type: Boolean

Required: No

Name

The service name.

Type: String

Required: No

Names

A collection of associated service names.

Type: Array of strings

Required: No

Type

The type associated to the service.

Type: String

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](#)

ErrorStatistics

Information about requests that failed with a 4xx Client Error status code.

Contents

OtherCount

The number of requests that failed with untracked 4xx Client Error status codes.

Type: Long

Required: No

ThrottleCount

The number of requests that failed with a 429 throttling status code.

Type: Long

Required: No

TotalCount

The total number of requests that failed with a 4xx Client Error status code.

Type: Long

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](#)

FaultRootCause

The root cause information for a trace summary fault.

Contents

ClientImpacting

A flag that denotes that the root cause impacts the trace client.

Type: Boolean

Required: No

Services

A list of corresponding services. A service identifies a segment and it contains a name, account ID, type, and inferred flag.

Type: Array of [FaultRootCauseService](#) 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](#)

FaultRootCauseEntity

A collection of segments and corresponding subsegments associated to a trace summary fault error.

Contents

Exceptions

The types and messages of the exceptions.

Type: Array of [RootCauseException](#) objects

Required: No

Name

The name of the entity.

Type: String

Required: No

Remote

A flag that denotes a remote subsegment.

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](#)

FaultRootCauseService

A collection of fields identifying the services in a trace summary fault.

Contents

AccountId

The account ID associated to the service.

Type: String

Required: No

EntityPath

The path of root cause entities found on the service.

Type: Array of [FaultRootCauseEntity](#) objects

Required: No

Inferred

A Boolean value indicating if the service is inferred from the trace.

Type: Boolean

Required: No

Name

The service name.

Type: String

Required: No

Names

A collection of associated service names.

Type: Array of strings

Required: No

Type

The type associated to the service.

Type: String

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](#)

FaultStatistics

Information about requests that failed with a 5xx Server Error status code.

Contents

OtherCount

The number of requests that failed with untracked 5xx Server Error status codes.

Type: Long

Required: No

TotalCount

The total number of requests that failed with a 5xx Server Error status code.

Type: Long

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](#)

ForecastStatistics

The predicted high and low fault count. This is used to determine if a service has become anomalous and if an insight should be created.

Contents

FaultCountHigh

The upper limit of fault counts for a service.

Type: Long

Required: No

FaultCountLow

The lower limit of fault counts for a service.

Type: Long

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](#)

GraphLink

The relation between two services.

Contents

DestinationTraceIds

Destination traces of a link relationship.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 35.

Required: No

ReferenceType

Relationship of a trace to the corresponding service.

Type: String

Required: No

SourceTraceId

Source trace of a link relationship.

Type: String

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](#)

Group

Details and metadata for a group.

Contents

FilterExpression

The filter expression defining the parameters to include traces.

Type: String

Required: No

GroupARN

The Amazon Resource Name (ARN) of the group generated based on the GroupName.

Type: String

Required: No

GroupName

The unique case-sensitive name of the group.

Type: String

Required: No

InsightsConfiguration

The structure containing configurations related to insights.

- The InsightsEnabled boolean can be set to true to enable insights for the group or false to disable insights for the group.
- The NotificationsEnabled boolean can be set to true to enable insights notifications through Amazon EventBridge for the group.

Type: [InsightsConfiguration](#) 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](#)

GroupSummary

Details for a group without metadata.

Contents

FilterExpression

The filter expression defining the parameters to include traces.

Type: String

Required: No

GroupARN

The ARN of the group generated based on the GroupName.

Type: String

Required: No

GroupName

The unique case-sensitive name of the group.

Type: String

Required: No

InsightsConfiguration

The structure containing configurations related to insights.

- The InsightsEnabled boolean can be set to true to enable insights for the group or false to disable insights for the group.
- The NotificationsEnabled boolean can be set to true to enable insights notifications. Notifications can only be enabled on a group with InsightsEnabled set to true.

Type: [InsightsConfiguration](#) 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](#)

HistogramEntry

An entry in a histogram for a statistic. A histogram maps the range of observed values on the X axis, and the prevalence of each value on the Y axis.

Contents

Count

The prevalence of the entry.

Type: Integer

Required: No

Value

The value of the entry.

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](#)

Http

Information about an HTTP request.

Contents

ClientIp

The IP address of the requestor.

Type: String

Required: No

HttpMethod

The request method.

Type: String

Required: No

HttpStatus

The response status.

Type: Integer

Required: No

HttpURL

The request URL.

Type: String

Required: No

UserAgent

The request's user agent string.

Type: String

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](#)

IndexingRule

Rule used to determine the server-side sampling rate for spans ingested through the CloudWatchLogs destination and indexed by X-Ray.

Contents

ModifiedAt

Displays when the rule was last modified, in Unix time seconds.

Type: Timestamp

Required: No

Name

The name of the indexing rule.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

Required: No

Rule

The indexing rule.

Type: [IndexingRuleValue](#) object

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

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](#)

IndexingRuleValue

The indexing rule configuration.

Contents

Important

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

Probabilistic

Indexing rule configuration that is used to probabilistically sample traceIds.

Type: [ProbabilisticRuleValue](#) 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](#)

IndexingRuleValueUpdate

Update to an indexing rule.

Contents

Important

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

Probabilistic

Indexing rule configuration that is used to probabilistically sample traceIds.

Type: [ProbabilisticRuleValueUpdate](#) 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](#)

Insight

When fault rates go outside of the expected range, X-Ray creates an insight. Insights tracks emergent issues within your applications.

Contents

Categories

The categories that label and describe the type of insight.

Type: Array of strings

Valid Values: FAULT

Required: No

ClientRequestImpactStatistics

The impact statistics of the client side service. This includes the number of requests to the client service and whether the requests were faults or okay.

Type: [RequestImpactStatistics](#) object

Required: No

EndTime

The time, in Unix seconds, at which the insight ended.

Type: Timestamp

Required: No

GroupARN

The Amazon Resource Name (ARN) of the group that the insight belongs to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 400.

Required: No

GroupName

The name of the group that the insight belongs to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

Required: No

InsightId

The insights unique identifier.

Type: String

Pattern: `[0-9a-fA-F]{8}-[0-9a-fA-F]{4}-[1-5][0-9a-fA-F]{3}-[89abAB][0-9a-fA-F]{3}-[0-9a-fA-F]{12}`

Required: No

RootCauseServiceId

Type: [ServiceId](#) object

Required: No

RootCauseServiceRequestImpactStatistics

The impact statistics of the root cause service. This includes the number of requests to the client service and whether the requests were faults or okay.

Type: [RequestImpactStatistics](#) object

Required: No

StartTime

The time, in Unix seconds, at which the insight began.

Type: Timestamp

Required: No

State

The current state of the insight.

Type: String

Valid Values: ACTIVE | CLOSED

Required: No

Summary

A brief description of the insight.

Type: String

Required: No

TopAnomalousServices

The service within the insight that is most impacted by the incident.

Type: Array of [AnomalousService](#) 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](#)

InsightEvent

X-Ray reevaluates insights periodically until they are resolved, and records each intermediate state in an event. You can review incident events in the Impact Timeline on the Inspect page in the X-Ray console.

Contents

ClientRequestImpactStatistics

The impact statistics of the client side service. This includes the number of requests to the client service and whether the requests were faults or okay.

Type: [RequestImpactStatistics](#) object

Required: No

EventTime

The time, in Unix seconds, at which the event was recorded.

Type: Timestamp

Required: No

RootCauseServiceRequestImpactStatistics

The impact statistics of the root cause service. This includes the number of requests to the client service and whether the requests were faults or okay.

Type: [RequestImpactStatistics](#) object

Required: No

Summary

A brief description of the event.

Type: String

Required: No

TopAnomalousServices

The service during the event that is most impacted by the incident.

Type: Array of [AnomalousService](#) 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](#)

InsightImpactGraphEdge

The connection between two service in an insight impact graph.

Contents

ReferenceId

Identifier of the edge. Unique within a service map.

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](#)

InsightImpactGraphService

Information about an application that processed requests, users that made requests, or downstream services, resources, and applications that an application used.

Contents

AccountId

Identifier of the AWS account in which the service runs.

Type: String

Required: No

Edges

Connections to downstream services.

Type: Array of [InsightImpactGraphEdge](#) objects

Required: No

Name

The canonical name of the service.

Type: String

Required: No

Names

A list of names for the service, including the canonical name.

Type: Array of strings

Required: No

ReferenceId

Identifier for the service. Unique within the service map.

Type: Integer

Required: No

Type

Identifier for the service. Unique within the service map.

- **AWS Resource** - The type of an AWS resource. For example, `AWS::EC2::Instance` for an application running on Amazon EC2 or `AWS::DynamoDB::Table` for an Amazon DynamoDB table that the application used.
- **AWS Service** - The type of an AWS service. For example, `AWS::DynamoDB` for downstream calls to Amazon DynamoDB that didn't target a specific table.
- **AWS Service** - The type of an AWS service. For example, `AWS::DynamoDB` for downstream calls to Amazon DynamoDB that didn't target a specific table.
- **remote** - A downstream service of indeterminate type.

Type: String

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](#)

InsightsConfiguration

The structure containing configurations related to insights.

Contents

InsightsEnabled

Set the InsightsEnabled value to true to enable insights or false to disable insights.

Type: Boolean

Required: No

NotificationsEnabled

Set the NotificationsEnabled value to true to enable insights notifications. Notifications can only be enabled on a group with InsightsEnabled set to true.

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](#)

InsightSummary

Information that describes an insight.

Contents

Categories

Categories The categories that label and describe the type of insight.

Type: Array of strings

Valid Values: FAULT

Required: No

ClientRequestImpactStatistics

The impact statistics of the client side service. This includes the number of requests to the client service and whether the requests were faults or okay.

Type: [RequestImpactStatistics](#) object

Required: No

EndTime

The time, in Unix seconds, at which the insight ended.

Type: Timestamp

Required: No

GroupARN

The Amazon Resource Name (ARN) of the group that the insight belongs to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 400.

Required: No

GroupName

The name of the group that the insight belongs to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

Required: No

InsightId

The insights unique identifier.

Type: String

Pattern: `[0-9a-fA-F]{8}-[0-9a-fA-F]{4}-[1-5][0-9a-fA-F]{3}-[89abAB][0-9a-fA-F]{3}-[0-9a-fA-F]{12}`

Required: No

LastUpdateTime

The time, in Unix seconds, that the insight was last updated.

Type: Timestamp

Required: No

RootCauseServiceId

Type: [ServiceId](#) object

Required: No

RootCauseServiceRequestImpactStatistics

The impact statistics of the root cause service. This includes the number of requests to the client service and whether the requests were faults or okay.

Type: [RequestImpactStatistics](#) object

Required: No

StartTime

The time, in Unix seconds, at which the insight began.

Type: Timestamp

Required: No

State

The current state of the insight.

Type: String

Valid Values: ACTIVE | CLOSED

Required: No

Summary

A brief description of the insight.

Type: String

Required: No

TopAnomalousServices

The service within the insight that is most impacted by the incident.

Type: Array of [AnomalousService](#) 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](#)

InstanceIdDetail

A list of EC2 instance IDs corresponding to the segments in a trace.

Contents

Id

The ID of a corresponding EC2 instance.

Type: String

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](#)

ProbabilisticRuleValue

The indexing rule configuration for probabilistic sampling.

Contents

DesiredSamplingPercentage

Configured sampling percentage of traceIds. Note that sampling can be subject to limits to ensure completeness of data.

Type: Double

Required: Yes

ActualSamplingPercentage

Applied sampling percentage of traceIds.

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](#)

ProbabilisticRuleValueUpdate

Update to the indexing rule configuration for probabilistic sampling.

Contents

DesiredSamplingPercentage

Configured sampling percentage of traceIds. Note that sampling can be subject to limits to ensure completeness of data.

Type: Double

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](#)

RequestImpactStatistics

Statistics that describe how the incident has impacted a service.

Contents

FaultCount

The number of requests that have resulted in a fault,

Type: Long

Required: No

OkCount

The number of successful requests.

Type: Long

Required: No

TotalCount

The total number of requests to the service.

Type: Long

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](#)

ResourceARNDetail

A list of resources ARNs corresponding to the segments in a trace.

Contents

ARN

The ARN of a corresponding resource.

Type: String

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](#)

ResourcePolicy

A resource policy grants one or more AWS services and accounts permissions to access X-Ray. Each resource policy is associated with a specific AWS account.

Contents

LastUpdatedTime

When the policy was last updated, in Unix time seconds.

Type: Timestamp

Required: No

PolicyDocument

The resource policy document, which can be up to 5kb in size.

Type: String

Required: No

PolicyName

The name of the resource policy. Must be unique within a specific AWS account.

Type: String

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

Pattern: `[\w+=, .@- ]+`

Required: No

PolicyRevisionId

Returns the current policy revision id for this policy name.

Type: String

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](#)

ResponseTimeRootCause

The root cause information for a response time warning.

Contents

ClientImpacting

A flag that denotes that the root cause impacts the trace client.

Type: Boolean

Required: No

Services

A list of corresponding services. A service identifies a segment and contains a name, account ID, type, and inferred flag.

Type: Array of [ResponseTimeRootCauseService](#) 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](#)

ResponseTimeRootCauseEntity

A collection of segments and corresponding subsegments associated to a response time warning.

Contents

Coverage

The type and messages of the exceptions.

Type: Double

Required: No

Name

The name of the entity.

Type: String

Required: No

Remote

A flag that denotes a remote subsegment.

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](#)

ResponseTimeRootCauseService

A collection of fields identifying the service in a response time warning.

Contents

AccountId

The account ID associated to the service.

Type: String

Required: No

EntityPath

The path of root cause entities found on the service.

Type: Array of [ResponseTimeRootCauseEntity](#) objects

Required: No

Inferred

A Boolean value indicating if the service is inferred from the trace.

Type: Boolean

Required: No

Name

The service name.

Type: String

Required: No

Names

A collection of associated service names.

Type: Array of strings

Required: No

Type

The type associated to the service.

Type: String

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](#)

RetrievedService

Retrieved information about an application that processed requests, users that made requests, or downstream services, resources, and applications that an application used.

Contents

Links

Relation between two 2 services.

Type: Array of [GraphLink](#) objects

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

Required: No

Service

Information about an application that processed requests, users that made requests, or downstream services, resources, and applications that an application used.

Type: [Service](#) 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](#)

RetrievedTrace

Retrieved collection of spans with matching trace IDs.

Contents

Duration

The length of time in seconds between the start time of the root span and the end time of the last span that completed.

Type: Double

Required: No

Id

The unique identifier for the span.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 35.

Required: No

Spans

Spans that comprise the trace.

Type: Array of [Span](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 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](#)

RootCauseException

The exception associated with a root cause.

Contents

Message

The message of the exception.

Type: String

Required: No

Name

The name of the exception.

Type: String

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](#)

SamplingRule

A sampling rule that services use to decide whether to instrument a request. Rule fields can match properties of the service, or properties of a request. The service can ignore rules that don't match its properties.

Contents

FixedRate

The percentage of matching requests to instrument, after the reservoir is exhausted.

Type: Double

Valid Range: Minimum value of 0. Maximum value of 1.

Required: Yes

Host

Matches the hostname from a request URL.

Type: String

Length Constraints: Maximum length of 64.

Required: Yes

HTTPMethod

Matches the HTTP method of a request.

Type: String

Length Constraints: Maximum length of 10.

Required: Yes

Priority

The priority of the sampling rule.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 9999.

Required: Yes

ReservoirSize

A fixed number of matching requests to instrument per second, prior to applying the fixed rate. The reservoir is not used directly by services, but applies to all services using the rule collectively.

Type: Integer

Valid Range: Minimum value of 0.

Required: Yes

ResourceARN

Matches the ARN of the AWS resource on which the service runs.

Type: String

Length Constraints: Maximum length of 500.

Required: Yes

ServiceName

Matches the name that the service uses to identify itself in segments.

Type: String

Length Constraints: Maximum length of 64.

Required: Yes

ServiceType

Matches the origin that the service uses to identify its type in segments.

Type: String

Length Constraints: Maximum length of 64.

Required: Yes

URLPath

Matches the path from a request URL.

Type: String

Length Constraints: Maximum length of 128.

Required: Yes

Version

The version of the sampling rule format (1).

Type: Integer

Valid Range: Minimum value of 1.

Required: Yes

Attributes

Matches attributes derived from the request.

Type: String to string map

Map Entries: Maximum number of 5 items.

Key Length Constraints: Minimum length of 1. Maximum length of 32.

Value Length Constraints: Minimum length of 1. Maximum length of 32.

Required: No

RuleARN

The ARN of the sampling rule. Specify a rule by either name or ARN, but not both.

Type: String

Required: No

RuleName

The name of the sampling rule. Specify a rule by either name or ARN, but not both.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

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](#)

SamplingRuleRecord

A [SamplingRule](#) and its metadata.

Contents

CreatedAt

When the rule was created, in Unix time seconds.

Type: Timestamp

Required: No

ModifiedAt

When the rule was last modified, in Unix time seconds.

Type: Timestamp

Required: No

SamplingRule

The sampling rule.

Type: [SamplingRule](#) 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](#)

SamplingRuleUpdate

A document specifying changes to a sampling rule's configuration.

Contents

Attributes

Matches attributes derived from the request.

Type: String to string map

Map Entries: Maximum number of 5 items.

Key Length Constraints: Minimum length of 1. Maximum length of 32.

Value Length Constraints: Minimum length of 1. Maximum length of 32.

Required: No

FixedRate

The percentage of matching requests to instrument, after the reservoir is exhausted.

Type: Double

Required: No

Host

Matches the hostname from a request URL.

Type: String

Length Constraints: Maximum length of 64.

Required: No

HTTPMethod

Matches the HTTP method of a request.

Type: String

Length Constraints: Maximum length of 10.

Required: No

Priority

The priority of the sampling rule.

Type: Integer

Required: No

ReservoirSize

A fixed number of matching requests to instrument per second, prior to applying the fixed rate. The reservoir is not used directly by services, but applies to all services using the rule collectively.

Type: Integer

Required: No

ResourceARN

Matches the ARN of the AWS resource on which the service runs.

Type: String

Length Constraints: Maximum length of 500.

Required: No

RuleARN

The ARN of the sampling rule. Specify a rule by either name or ARN, but not both.

Type: String

Required: No

RuleName

The name of the sampling rule. Specify a rule by either name or ARN, but not both.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

Required: No

ServiceName

Matches the name that the service uses to identify itself in segments.

Type: String

Length Constraints: Maximum length of 64.

Required: No

ServiceType

Matches the origin that the service uses to identify its type in segments.

Type: String

Length Constraints: Maximum length of 64.

Required: No

URLPath

Matches the path from a request URL.

Type: String

Length Constraints: Maximum length of 128.

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](#)

SamplingStatisticsDocument

Request sampling results for a single rule from a service. Results are for the last 10 seconds unless the service has been assigned a longer reporting interval after a previous call to [GetSamplingTargets](#).

Contents

ClientID

A unique identifier for the service in hexadecimal.

Type: String

Length Constraints: Fixed length of 24.

Required: Yes

RequestCount

The number of requests that matched the rule.

Type: Integer

Valid Range: Minimum value of 0.

Required: Yes

RuleName

The name of the sampling rule.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

Required: Yes

SampledCount

The number of requests recorded.

Type: Integer

Valid Range: Minimum value of 0.

Required: Yes

Timestamp

The current time, in [ISO-8601 format](#) (YYYY-MM-DDThh:mm:ss).

Type: Timestamp

Required: Yes

BorrowCount

The number of requests recorded with borrowed reservoir quota.

Type: Integer

Valid Range: Minimum value of 0.

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](#)

SamplingStatisticSummary

Aggregated request sampling data for a sampling rule across all services for a 10-second window.

Contents

BorrowCount

The number of requests recorded with borrowed reservoir quota.

Type: Integer

Required: No

RequestCount

The number of requests that matched the rule.

Type: Integer

Required: No

RuleName

The name of the sampling rule.

Type: String

Required: No

SampledCount

The number of requests recorded.

Type: Integer

Required: No

Timestamp

The start time of the reporting window, in Unix time seconds.

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](#)

SamplingStrategy

The name and value of a sampling rule to apply to a trace summary.

Contents

Name

The name of a sampling rule.

Type: String

Valid Values: `PartialScan` | `FixedRate`

Required: No

Value

The value of a sampling rule.

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](#)

SamplingTargetDocument

Temporary changes to a sampling rule configuration. To meet the global sampling target for a rule, X-Ray calculates a new reservoir for each service based on the recent sampling results of all services that called [GetSamplingTargets](#).

Contents

FixedRate

The percentage of matching requests to instrument, after the reservoir is exhausted.

Type: Double

Required: No

Interval

The number of seconds for the service to wait before getting sampling targets again.

Type: Integer

Required: No

ReservoirQuota

The number of requests per second that X-Ray allocated for this service.

Type: Integer

Required: No

ReservoirQuotaTTL

When the reservoir quota expires.

Type: Timestamp

Required: No

RuleName

The name of the sampling rule.

Type: String

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](#)

Segment

A segment from a trace that has been ingested by the X-Ray service. The segment can be compiled from documents uploaded with [PutTraceSegments](#), or an inferred segment for a downstream service, generated from a subsegment sent by the service that called it.

For the full segment document schema, see [AWS X-Ray segment documents](#) in the *AWS X-Ray Developer Guide*.

Contents

Document

The segment document.

Type: String

Length Constraints: Minimum length of 1.

Required: No

Id

The segment's ID.

Type: String

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](#)

Service

Information about an application that processed requests, users that made requests, or downstream services, resources, and applications that an application used.

Contents

AccountId

Identifier of the AWS account in which the service runs.

Type: String

Required: No

DurationHistogram

A histogram that maps the spread of service durations.

Type: Array of [HistogramEntry](#) objects

Required: No

Edges

Connections to downstream services.

Type: Array of [Edge](#) objects

Required: No

EndTime

The end time of the last segment that the service generated.

Type: Timestamp

Required: No

Name

The canonical name of the service.

Type: String

Required: No

Names

A list of names for the service, including the canonical name.

Type: Array of strings

Required: No

ReferenceId

Identifier for the service. Unique within the service map.

Type: Integer

Required: No

ResponseTimeHistogram

A histogram that maps the spread of service response times.

Type: Array of [HistogramEntry](#) objects

Required: No

Root

Indicates that the service was the first service to process a request.

Type: Boolean

Required: No

StartTime

The start time of the first segment that the service generated.

Type: Timestamp

Required: No

State

The service's state.

Type: String

Required: No

SummaryStatistics

Aggregated statistics for the service.

Type: [ServiceStatistics](#) object

Required: No

Type

The type of service.

- **AWS Resource** - The type of an AWS resource. For example, `AWS::EC2::Instance` for an application running on Amazon EC2 or `AWS::DynamoDB::Table` for an Amazon DynamoDB table that the application used.
- **AWS Service** - The type of an AWS service. For example, `AWS::DynamoDB` for downstream calls to Amazon DynamoDB that didn't target a specific table.
- `client` - Represents the clients that sent requests to a root service.
- `remote` - A downstream service of indeterminate type.

Type: String

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](#)

ServiceId

Contents

AccountId

Type: String

Required: No

Name

Type: String

Required: No

Names

Type: Array of strings

Required: No

Type

Type: String

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](#)

ServiceStatistics

Response statistics for a service.

Contents

ErrorStatistics

Information about requests that failed with a 4xx Client Error status code.

Type: [ErrorStatistics](#) object

Required: No

FaultStatistics

Information about requests that failed with a 5xx Server Error status code.

Type: [FaultStatistics](#) object

Required: No

OkCount

The number of requests that completed with a 2xx Success status code.

Type: Long

Required: No

TotalCount

The total number of completed requests.

Type: Long

Required: No

TotalResponseTime

The aggregate response time of completed requests.

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](#)

Span

A span from a trace that has been ingested by the X-Ray service. A span represents a unit of work or an operation performed by a service.

Contents

Document

The span document.

Type: String

Length Constraints: Maximum length of 204800.

Required: No

Id

The span ID.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 16.

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 map that contains tag keys and tag values to attach to an AWS X-Ray group or sampling rule. For more information about ways to use tags, see [Tagging AWS resources](#) in the *AWS General Reference*.

The following restrictions apply to tags:

- Maximum number of user-applied tags per resource: 50
- Tag keys and values are case sensitive.
- Don't use `aws :` as a prefix for keys; it's reserved for AWS use. You cannot edit or delete system tags.

Contents

Key

A tag key, such as `Stage` or `Name`. A tag key cannot be empty. The key can be a maximum of 128 characters, and can contain only Unicode letters, numbers, or separators, or the following special characters: `+ - = . _ : /`

Type: String

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

Required: Yes

Value

An optional tag value, such as `Production` or `test-only`. The value can be a maximum of 255 characters, and contain only Unicode letters, numbers, or separators, or the following special characters: `+ - = . _ : /`

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](#)

TelemetryRecord

Contents

Timestamp

Type: Timestamp

Required: Yes

BackendConnectionErrors

Type: [BackendConnectionErrors](#) object

Required: No

SegmentsReceivedCount

Type: Integer

Required: No

SegmentsRejectedCount

Type: Integer

Required: No

SegmentsSentCount

Type: Integer

Required: No

SegmentsSpilloverCount

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](#)

TimeSeriesServiceStatistics

A list of TimeSeriesStatistic structures.

Contents

EdgeSummaryStatistics

Response statistics for an edge.

Type: [EdgeStatistics](#) object

Required: No

ResponseTimeHistogram

The response time histogram for the selected entities.

Type: Array of [HistogramEntry](#) objects

Required: No

ServiceForecastStatistics

The forecasted high and low fault count values.

Type: [ForecastStatistics](#) object

Required: No

ServiceSummaryStatistics

Response statistics for a service.

Type: [ServiceStatistics](#) object

Required: No

Timestamp

Timestamp of the window for which statistics are aggregated.

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](#)

Trace

A collection of segment documents with matching trace IDs.

Contents

Duration

The length of time in seconds between the start time of the earliest segment that started and the end time of the last segment that completed.

Type: Double

Required: No

Id

The unique identifier for the request that generated the trace's segments and subsegments.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 35.

Required: No

LimitExceeded

LimitExceeded is set to true when the trace has exceeded the `Trace` document size limit. For more information about this limit and other X-Ray limits and quotas, see [AWS X-Ray endpoints and quotas](#).

Type: Boolean

Required: No

Segments

Segment documents for the segments and subsegments that comprise the trace.

Type: Array of [Segment](#) 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](#)

TraceSummary

Metadata generated from the segment documents in a trace.

Contents

Annotations

Annotations from the trace's segment documents.

Type: String to array of [ValueWithServiceIds](#) objects map

Required: No

AvailabilityZones

A list of Availability Zones for any zone corresponding to the trace segments.

Type: Array of [AvailabilityZoneDetail](#) objects

Required: No

Duration

The length of time in seconds between the start time of the earliest segment that started and the end time of the last segment that completed.

Type: Double

Required: No

EntryPoint

The root of a trace.

Type: [ServiceId](#) object

Required: No

ErrorRootCauses

A collection of ErrorRootCause structures corresponding to the trace segments.

Type: Array of [ErrorRootCause](#) objects

Required: No

FaultRootCauses

A collection of FaultRootCause structures corresponding to the trace segments.

Type: Array of [FaultRootCause](#) objects

Required: No

HasError

The root segment document has a 400 series error.

Type: Boolean

Required: No

HasFault

The root segment document has a 500 series error.

Type: Boolean

Required: No

HasThrottle

One or more of the segment documents has a 429 throttling error.

Type: Boolean

Required: No

Http

Information about the HTTP request served by the trace.

Type: [Http](#) object

Required: No

Id

The unique identifier for the request that generated the trace's segments and subsegments.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 35.

Required: No

InstanceIds

A list of EC2 instance IDs for any instance corresponding to the trace segments.

Type: Array of [InstanceDetail](#) objects

Required: No

IsPartial

One or more of the segment documents is in progress.

Type: Boolean

Required: No

MatchedEventTime

The matched time stamp of a defined event.

Type: Timestamp

Required: No

ResourceARNs

A list of resource ARNs for any resource corresponding to the trace segments.

Type: Array of [ResourceARNDetail](#) objects

Required: No

ResponseTime

The length of time in seconds between the start and end times of the root segment. If the service performs work asynchronously, the response time measures the time before the response is sent to the user, while the duration measures the amount of time before the last traced activity completes.

Type: Double

Required: No

ResponseTimeRootCauses

A collection of ResponseTimeRootCause structures corresponding to the trace segments.

Type: Array of [ResponseTimeRootCause](#) objects

Required: No

Revision

The revision number of a trace.

Type: Integer

Required: No

ServiceIds

Service IDs from the trace's segment documents.

Type: Array of [ServiceId](#) objects

Required: No

StartTime

The start time of a trace, based on the earliest trace segment start time.

Type: Timestamp

Required: No

Users

Users from the trace's segment documents.

Type: Array of [TraceUser](#) 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](#)

TraceUser

Information about a user recorded in segment documents.

Contents

ServiceIds

Services that the user's request hit.

Type: Array of [ServiceId](#) objects

Required: No

UserName

The user's name.

Type: String

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](#)

UnprocessedStatistics

Sampling statistics from a call to [GetSamplingTargets](#) that X-Ray could not process.

Contents

ErrorCode

The error code.

Type: String

Required: No

Message

The error message.

Type: String

Required: No

RuleName

The name of the sampling rule.

Type: String

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](#)

UnprocessedTraceSegment

Information about a segment that failed processing.

Contents

ErrorCode

The error that caused processing to fail.

Type: String

Required: No

Id

The segment's ID.

Type: String

Required: No

Message

The error message.

Type: String

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](#)

ValueWithServiceIds

Information about a segment annotation.

Contents

AnnotationValue

Values of the annotation.

Type: [AnnotationValue](#) object

Required: No

ServiceIds

Services to which the annotation applies.

Type: Array of [ServiceId](#) 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](#)

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