API Reference

Amazon ElastiCache

API Version 2015-02-02

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

Amazon ElastiCache: 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**
AddTagsToResource 5
Request Parameters 5
Response Elements 6
Errors 6
Examples 8
See Also 8
AuthorizeCacheSecurityGroupIngress 10
Request Parameters 10
Response Elements 11
Errors 11
Examples 12
See Also 13
BatchApplyUpdateAction 14
Request Parameters 14
Response Elements 14
Errors 15
See Also 15
BatchStopUpdateAction 17
Request Parameters 17
Response Elements 17
Errors 18
See Also 18
CompleteMigration 20
Request Parameters 20
Response Elements 20
Errors 20
See Also 21
CopyServerlessCacheSnapshot 22
Request Parameters 22
Response Elements 23
Errors 23
Examples 24

See Also ... 25
CopySnapshot ... 26
Request Parameters ... 27
Response Elements ... 28
Errors ... 28
Examples ... 30
See Also ... 30
CreateCacheCluster ... 32
Request Parameters ... 32
Response Elements ... 43
Errors ... 43
Examples ... 45
See Also ... 46
CreateCacheParameterGroup ... 48
Request Parameters ... 48
Response Elements ... 49
Errors ... 49
Examples ... 50
See Also ... 51
CreateCacheSecurityGroup ... 52
Request Parameters ... 52
Response Elements ... 53
Errors ... 53
Examples ... 54
See Also ... 55
CreateCacheSubnetGroup ... 56
Request Parameters ... 56
Response Elements ... 57
Errors ... 57
Examples ... 58
See Also ... 59
CreateGlobalReplicationGroup ... 61
Request Parameters ... 61
Response Elements ... 62
Errors ... 62
See Also ... 63

CreateReplicationGroup ... 64
Request Parameters ... 65
Response Elements ... 79
Errors ... 79
Examples ... 82
See Also ... 85
CreateServerlessCache ... 86
Request Parameters ... 86
Response Elements ... 88
Errors ... 88
Examples ... 90
See Also ... 93
CreateServerlessCacheSnapshot ... 94
Request Parameters ... 94
Response Elements ... 95
Errors ... 95
Examples ... 96
See Also ... 97
CreateSnapshot ... 98
Request Parameters ... 98
Response Elements ... 99
Errors ... 99
Examples ... 101
See Also ... 102
CreateUser ... 103
Request Parameters ... 103
Response Elements ... 104
Errors ... 106
See Also ... 107
CreateUserGroup ... 108
Request Parameters ... 108
Response Elements ... 109
Errors ... 110
See Also ... 111
DecreaseNodeGroupsInGlobalReplicationGroup ... 113
Request Parameters ... 113

Response Elements ... 114
Errors ... 114
See Also ... 115
DecreaseReplicaCount ... 116
Request Parameters ... 116
Response Elements ... 117
Errors ... 117
Examples ... 119
See Also ... 120
DeleteCacheCluster ... 122
Request Parameters ... 122
Response Elements ... 123
Errors ... 123
Examples ... 124
See Also ... 125
DeleteCacheParameterGroup ... 127
Request Parameters ... 127
Errors ... 127
Examples ... 128
See Also ... 128
DeleteCacheSecurityGroup ... 130
Request Parameters ... 130
Errors ... 130
Examples ... 131
See Also ... 132
DeleteCacheSubnetGroup ... 133
Request Parameters ... 133
Errors ... 133
Examples ... 134
See Also ... 134
DeleteGlobalReplicationGroup ... 136
Request Parameters ... 136
Response Elements ... 136
Errors ... 137
See Also ... 137
DeleteReplicationGroup ... 139

Request Parameters ... 139
Response Elements ... 140
Errors ... 140
Examples ... 141
See Also ... 142
DeleteServerlessCache ... 143
Request Parameters ... 143
Response Elements ... 143
Errors ... 144
Examples ... 145
See Also ... 146
DeleteServerlessCacheSnapshot ... 147
Request Parameters ... 147
Response Elements ... 147
Errors ... 147
See Also ... 148
DeleteSnapshot ... 149
Request Parameters ... 149
Response Elements ... 149
Errors ... 149
Examples ... 150
See Also ... 151
DeleteUser ... 153
Request Parameters ... 153
Response Elements ... 153
Errors ... 154
See Also ... 155
DeleteUserGroup ... 157
Request Parameters ... 157
Response Elements ... 157
Errors ... 158
See Also ... 159
DescribeCacheClusters ... 160
Request Parameters ... 160
Response Elements ... 161
Errors ... 162

Examples ... 162
See Also ... 164
DescribeCacheEngineVersions ... 165
Request Parameters ... 165
Response Elements ... 166
Errors ... 167
Examples ... 167
See Also ... 168
DescribeCacheParameterGroups ... 169
Request Parameters ... 169
Response Elements ... 169
Errors ... 170
Examples ... 170
See Also ... 172
DescribeCacheParameters ... 173
Request Parameters ... 173
Response Elements ... 174
Errors ... 174
Examples ... 175
See Also ... 176
DescribeCacheSecurityGroups ... 178
Request Parameters ... 178
Response Elements ... 179
Errors ... 179
Examples ... 179
See Also ... 180
DescribeCacheSubnetGroups ... 182
Request Parameters ... 182
Response Elements ... 183
Errors ... 183
Examples ... 183
See Also ... 185
DescribeEngineDefaultParameters ... 186
Request Parameters ... 186
Response Elements ... 187
Errors ... 187

Examples ... 187
See Also ... 189
DescribeEvents ... 190
Request Parameters ... 190
Response Elements ... 192
Errors ... 192
Examples ... 192
See Also ... 193
DescribeGlobalReplicationGroups ... 195
Request Parameters ... 195
Response Elements ... 196
Errors ... 196
See Also ... 196
DescribeReplicationGroups ... 198
Request Parameters ... 198
Response Elements ... 199
Errors ... 199
Examples ... 200
See Also ... 201
DescribeReservedCacheNodes ... 203
Request Parameters ... 203
Response Elements ... 207
Errors ... 207
Examples ... 208
See Also ... 209
DescribeReservedCacheNodesOfferings ... 210
Request Parameters ... 210
Response Elements ... 214
Errors ... 214
Examples ... 215
See Also ... 216
DescribeServerlessCaches ... 217
Request Parameters ... 217
Response Elements ... 217
Errors ... 218
Examples ... 218

See Also ... 220
DescribeServerlessCacheSnapshots ... 221
Request Parameters ... 221
Response Elements ... 222
Errors ... 222
See Also ... 223
DescribeServiceUpdates ... 225
Request Parameters ... 225
Response Elements ... 226
Errors ... 226
See Also ... 226
DescribeSnapshots ... 228
Request Parameters ... 228
Response Elements ... 229
Errors ... 230
Examples ... 231
See Also ... 232
DescribeUpdateActions ... 233
Request Parameters ... 233
Response Elements ... 235
Errors ... 235
See Also ... 236
DescribeUserGroups ... 237
Request Parameters ... 237
Response Elements ... 237
Errors ... 238
See Also ... 238
DescribeUsers ... 240
Request Parameters ... 240
Response Elements ... 241
Errors ... 241
See Also ... 242
DisassociateGlobalReplicationGroup ... 243
Request Parameters ... 243
Response Elements ... 243
Errors ... 244

See Also ... 244
ExportServerlessCacheSnapshot ... 246
Request Parameters ... 246
Response Elements ... 246
Errors ... 246
See Also ... 247
FailoverGlobalReplicationGroup ... 249
Request Parameters ... 249
Response Elements ... 249
Errors ... 250
See Also ... 250
IncreaseNodeGroupsInGlobalReplicationGroup ... 252
Request Parameters ... 252
Response Elements ... 252
Errors ... 253
See Also ... 253
IncreaseReplicaCount ... 255
Request Parameters ... 255
Response Elements ... 256
Errors ... 256
Examples ... 258
See Also ... 259
ListAllowedNodeTypeModifications ... 260
Request Parameters ... 260
Response Elements ... 261
Errors ... 261
Examples ... 262
See Also ... 263
ListTagsForResource ... 264
Request Parameters ... 264
Response Elements ... 264
Errors ... 264
Examples ... 266
See Also ... 267
ModifyCacheCluster ... 268
Request Parameters ... 268

Response Elements 276
Errors 277
Examples 278
See Also 280
ModifyCacheParameterGroup 281
Request Parameters 281
Response Elements 281
Errors 281
Examples 282
See Also 283
ModifyCacheSubnetGroup 284
Request Parameters 284
Response Elements 284
Errors 285
Examples 286
See Also 287
ModifyGlobalReplicationGroup 288
Request Parameters 288
Response Elements 289
Errors 290
See Also 290
ModifyReplicationGroup 292
Request Parameters 292
Response Elements 300
Errors 300
Examples 302
See Also 304
ModifyReplicationGroupShardConfiguration 306
Request Parameters 306
Response Elements 308
Errors 308
Examples 309
See Also 311
ModifyServerlessCache 313
Request Parameters 313
Response Elements 315

Errors ... 315
Examples ... 316
See Also ... 317
ModifyUser ... 319
Request Parameters ... 319
Response Elements ... 320
Errors ... 322
See Also ... 322
ModifyUserGroup ... 324
Request Parameters ... 324
Response Elements ... 325
Errors ... 326
See Also ... 327
PurchaseReservedCacheNodesOffering ... 329
Request Parameters ... 329
Response Elements ... 330
Errors ... 330
Examples ... 331
See Also ... 332
RebalanceSlotsInGlobalReplicationGroup ... 333
Request Parameters ... 333
Response Elements ... 333
Errors ... 333
See Also ... 334
RebootCacheCluster ... 335
Request Parameters ... 335
Response Elements ... 335
Errors ... 336
Examples ... 336
See Also ... 338
RemoveTagsFromResource ... 339
Request Parameters ... 339
Response Elements ... 339
Errors ... 340
Examples ... 342
See Also ... 342

ResetCacheParameterGroup 344
Request Parameters 344
Response Elements 344
Errors 345
Examples 346
See Also 346
RevokeCacheSecurityGroupIngress 348
Request Parameters 348
Response Elements 348
Errors 349
Examples 349
See Also 350
StartMigration 352
Request Parameters 352
Response Elements 352
Errors 352
See Also 353
TestFailover 354
Request Parameters 355
Response Elements 355
Errors 355
Examples 357
See Also 357
TestMigration 359
Request Parameters 359
Response Elements 359
Errors 359
See Also 360
Data Types 361
Authentication 364
Contents 364
See Also 364
AuthenticationMode 365
Contents 365
See Also 365
AvailabilityZone 367

Contents ... 367
See Also ... 367
CacheCluster ... 368
Contents ... 368
See Also ... 377
CacheEngineVersion ... 378
Contents ... 378
See Also ... 379
CacheNode ... 380
Contents ... 382
See Also ... 384
CacheNodeTypeSpecificParameter ... 385
Contents ... 385
See Also ... 387
CacheNodeTypeSpecificValue ... 388
Contents ... 388
See Also ... 388
CacheNodeUpdateStatus ... 389
Contents ... 389
See Also ... 390
CacheParameterGroup ... 392
Contents ... 392
See Also ... 393
CacheParameterGroupStatus ... 394
Contents ... 394
See Also ... 394
CacheSecurityGroup ... 396
Contents ... 396
See Also ... 397
CacheSecurityGroupMembership ... 398
Contents ... 398
See Also ... 398
CacheSubnetGroup ... 399
Contents ... 399
See Also ... 400
CacheUsageLimits ... 401

Contents ... 401
See Also ... 401
CloudWatchLogsDestinationDetails ... 402
Contents ... 402
See Also ... 402
ConfigureShard ... 403
Contents ... 403
See Also ... 404
CustomerNodeEndpoint ... 405
Contents ... 405
See Also ... 405
DataStorage ... 406
Contents ... 406
See Also ... 406
DestinationDetails ... 408
Contents ... 408
See Also ... 408
EC2SecurityGroup ... 409
Contents ... 409
See Also ... 409
ECPUPerSecond ... 411
Contents ... 411
See Also ... 411
Endpoint ... 412
Contents ... 412
See Also ... 412
EngineDefaults ... 413
Contents ... 413
See Also ... 414
Event ... 415
Contents ... 415
See Also ... 416
Filter ... 417
Contents ... 417
See Also ... 417
GlobalNodeGroup ... 419

Contents ... 419
See Also ... 419
GlobalReplicationGroup ... 420
Contents ... 420
See Also ... 423
GlobalReplicationGroupInfo ... 424
Contents ... 424
See Also ... 424
GlobalReplicationGroupMember ... 425
Contents ... 425
See Also ... 426
KinesisFirehoseDestinationDetails ... 427
Contents ... 427
See Also ... 427
LogDeliveryConfiguration ... 428
Contents ... 428
See Also ... 429
LogDeliveryConfigurationRequest ... 430
Contents ... 430
See Also ... 431
NodeGroup ... 432
Contents ... 432
See Also ... 433
NodeGroupConfiguration ... 434
Contents ... 434
See Also ... 435
NodeGroupMember ... 436
Contents ... 436
See Also ... 437
NodeGroupMemberUpdateStatus ... 438
Contents ... 438
See Also ... 439
NodeGroupUpdateStatus ... 441
Contents ... 441
See Also ... 441
NodeSnapshot ... 442

Contents ... 442
See Also ... 443
NotificationConfiguration ... 444
Contents ... 444
See Also ... 444
Parameter ... 445
Contents ... 445
See Also ... 446
ParameterNameValue ... 448
Contents ... 448
See Also ... 448
PendingLogDeliveryConfiguration ... 449
Contents ... 449
See Also ... 450
PendingModifiedValues ... 451
Contents ... 451
See Also ... 453
ProcessedUpdateAction ... 454
Contents ... 454
See Also ... 455
RecurringCharge ... 456
Contents ... 456
See Also ... 456
RegionalConfiguration ... 457
Contents ... 457
See Also ... 457
ReplicationGroup ... 459
Contents ... 459
See Also ... 466
ReplicationGroupPendingModifiedValues ... 467
Contents ... 467
See Also ... 469
ReservedCacheNode ... 470
Contents ... 470
See Also ... 474
ReservedCacheNodesOffering ... 476

Contents 476
See Also 479
ReshardingConfiguration 481
Contents 481
See Also 481
ReshardingStatus 483
Contents 483
See Also 483
ScaleConfig 484
Contents 484
See Also 484
SecurityGroupMembership 485
Contents 485
See Also 485
ServerlessCache 486
Contents 486
See Also 489
ServerlessCacheConfiguration 490
Contents 490
See Also 490
ServerlessCacheSnapshot 492
Contents 492
See Also 494
ServiceUpdate 495
Contents 495
See Also 497
SlotMigration 498
Contents 498
See Also 498
Snapshot 499
Contents 499
See Also 507
Subnet 508
Contents 508
See Also 509
SubnetOutpost 510

Contents 510
See Also 510
Tag 511
Contents 511
See Also 511
TimeRangeFilter 512
Contents 512
See Also 512
UnprocessedUpdateAction 513
Contents 513
See Also 514
UpdateAction 515
Contents 515
See Also 518
User 519
Contents 519
See Also 520
UserGroup 522
Contents 522
See Also 523
UserGroupPendingChanges 525
Contents 525
See Also 525
UserGroupsUpdateStatus 527
Contents 527
See Also 527
Common Parameters 529
Common Errors 532

Welcome

Amazon ElastiCache is a web service that makes it easier to set up, operate, and scale a distributed cache in the cloud.

With ElastiCache, customers get all of the benefits of a high-performance, in-memory cache with less of the administrative burden involved in launching and managing a distributed cache. The service makes setup, scaling, and cluster failure handling much simpler than in a self-managed cache deployment.

In addition, through integration with Amazon CloudWatch, customers get enhanced visibility into the key performance statistics associated with their cache and can receive alarms if a part of their cache runs hot.

This document was last published on September 4, 2025.

Actions

The following actions are supported:

- AddTagsToResource
- AuthorizeCacheSecurityGroupIngress
- BatchApplyUpdateAction
- BatchStopUpdateAction
- CompleteMigration
- CopyServerlessCacheSnapshot
- CopySnapshot
- CreateCacheCluster
- CreateCacheParameterGroup
- CreateCacheSecurityGroup
- CreateCacheSubnetGroup
- CreateGlobalReplicationGroup
- CreateReplicationGroup
- CreateServerlessCache
- CreateServerlessCacheSnapshot
- CreateSnapshot
- CreateUser
- CreateUserGroup
- DecreaseNodeGroupsInGlobalReplicationGroup
- DecreaseReplicaCount
- DeleteCacheCluster
- DeleteCacheParameterGroup
- DeleteCacheSecurityGroup
- DeleteCacheSubnetGroup
- DeleteGlobalReplicationGroup
- DeleteReplicationGroup
- DeleteServerlessCache

- DeleteServerlessCacheSnapshot
- DeleteSnapshot
- DeleteUser
- DeleteUserGroup
- DescribeCacheClusters
- DescribeCacheEngineVersions
- DescribeCacheParameterGroups
- DescribeCacheParameters
- DescribeCacheSecurityGroups
- DescribeCacheSubnetGroups
- DescribeEngineDefaultParameters
- DescribeEvents
- DescribeGlobalReplicationGroups
- DescribeReplicationGroups
- DescribeReservedCacheNodes
- DescribeReservedCacheNodesOfferings
- DescribeServerlessCaches
- DescribeServerlessCacheSnapshots
- DescribeServiceUpdates
- DescribeSnapshots
- DescribeUpdateActions
- DescribeUserGroups
- DescribeUsers
- DisassociateGlobalReplicationGroup
- ExportServerlessCacheSnapshot
- FailoverGlobalReplicationGroup
- IncreaseNodeGroupsInGlobalReplicationGroup
- IncreaseReplicaCount
- ListAllowedNodeTypeModifications
- ListTagsForResource

- ModifyCacheCluster
- ModifyCacheParameterGroup
- ModifyCacheSubnetGroup
- ModifyGlobalReplicationGroup
- ModifyReplicationGroup
- ModifyReplicationGroupShardConfiguration
- ModifyServerlessCache
- ModifyUser
- ModifyUserGroup
- PurchaseReservedCacheNodesOffering
- RebalanceSlotsInGlobalReplicationGroup
- RebootCacheCluster
- RemoveTagsFromResource
- ResetCacheParameterGroup
- RevokeCacheSecurityGroupIngress
- StartMigration
- TestFailover
- TestMigration

AddTagsToResource

A tag is a key-value pair where the key and value are case-sensitive. You can use tags to categorize and track all your ElastiCache resources, with the exception of global replication group. When you add or remove tags on replication groups, those actions will be replicated to all nodes in the replication group. For more information, see Resource-level permissions.

For example, you can use cost-allocation tags to your ElastiCache resources, Amazon generates a cost allocation report as a comma-separated value (CSV) file with your usage and costs aggregated by your tags. You can apply tags that represent business categories (such as cost centers, application names, or owners) to organize your costs across multiple services.

For more information, see Using Cost Allocation Tags in Amazon ElastiCache in the *ElastiCache User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ResourceName

The Amazon Resource Name (ARN) of the resource to which the tags are to be added, for example `arn:aws:elasticache:us-west-2:0123456789:cluster:myCluster` or `arn:aws:elasticache:us-west-2:0123456789:snapshot:mySnapshot`. ElastiCache resources are *cluster* and *snapshot*.

For more information about ARNs, see Amazon Resource Names (ARNs) and Amazon Service Namespaces.

Type: String

Required: Yes

Tags.Tag.N

A list of tags to be added to this resource. A tag is a key-value pair. A tag key must be accompanied by a tag value, although null is accepted.

Type: Array of Tag objects

Required: Yes

Response Elements

The following element is returned by the service.

TagList.Tag.N

A list of tags as key-value pairs.

Type: Array of Tag objects

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterNotFound

The requested cluster ID does not refer to an existing cluster.

HTTP Status Code: 404

CacheParameterGroupNotFound

The requested cache parameter group name does not refer to an existing cache parameter group.

HTTP Status Code: 404

CacheSecurityGroupNotFound

The requested cache security group name does not refer to an existing cache security group.

HTTP Status Code: 404

CacheSubnetGroupNotFoundFault

The requested cache subnet group name does not refer to an existing cache subnet group.

HTTP Status Code: 400

InvalidARN

The requested Amazon Resource Name (ARN) does not refer to an existing resource.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

InvalidServerlessCacheSnapshotStateFault

The state of the serverless cache snapshot was not received. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 400

InvalidServerlessCacheStateFault

The account for these credentials is not currently active.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

ReservedCacheNodeNotFound

The requested reserved cache node was not found.

HTTP Status Code: 404

ServerlessCacheNotFoundFault

The serverless cache was not found or does not exist.

HTTP Status Code: 404

ServerlessCacheSnapshotNotFoundFault

This serverless cache snapshot could not be found or does not exist. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 404

SnapshotNotFoundFault

The requested snapshot name does not refer to an existing snapshot.

HTTP Status Code: 404

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

UserGroupNotFound

The user group was not found or does not exist

HTTP Status Code: 404

UserNotFound

The user does not exist or could not be found.

HTTP Status Code: 404

Examples

AddTagsToResource

This example illustrates one usage of AddTagsToResource.

Sample Request

```
https://elasticache.us-east-1.amazonaws.com/
   ?Action=AddTagsToResource
   &ResourceName=arn:aws:elasticache:us-west-2:0123456789:cluster:myCluster
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Tags.Tag.1.Key=Service
   &Tags.Tag.1.Value=elasticache
   &Tags.Tag.2.Key=Region
   &Tags.Tag.2.Value=us-west-2
   &Version=2015-02-02
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

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

AuthorizeCacheSecurityGroupIngress

Allows network ingress to a cache security group. Applications using ElastiCache must be running on Amazon EC2, and Amazon EC2 security groups are used as the authorization mechanism.

> **Note**
>
> You cannot authorize ingress from an Amazon EC2 security group in one region to an ElastiCache cluster in another region.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheSecurityGroupName

The cache security group that allows network ingress.

Type: String

Required: Yes

EC2SecurityGroupName

The Amazon EC2 security group to be authorized for ingress to the cache security group.

Type: String

Required: Yes

EC2SecurityGroupOwnerId

The Amazon account number of the Amazon EC2 security group owner. Note that this is not the same thing as an Amazon access key ID - you must provide a valid Amazon account number for this parameter.

Type: String

Required: Yes

Response Elements

The following element is returned by the service.

CacheSecurityGroup

Represents the output of one of the following operations:

- `AuthorizeCacheSecurityGroupIngress`
- `CreateCacheSecurityGroup`
- `RevokeCacheSecurityGroupIngress`

Type: CacheSecurityGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

AuthorizationAlreadyExists

The specified Amazon EC2 security group is already authorized for the specified cache security group.

HTTP Status Code: 400

CacheSecurityGroupNotFound

The requested cache security group name does not refer to an existing cache security group.

HTTP Status Code: 404

InvalidCacheSecurityGroupState

The current state of the cache security group does not allow deletion.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

Examples

AuthorizeCacheSecurityGroupIngress

This example illustrates one usage of AuthorizeCacheSecurityGroupIngress.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=AuthorizeCacheSecurityGroupIngress
   &EC2SecurityGroupName=default
   &CacheSecurityGroupName=mygroup
   &EC2SecurityGroupOwnerId=1234-5678-1234
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Example

This example illustrates one usage of AuthorizeCacheSecurityGroupIngress.

Sample Response

```
<AuthorizeCacheSecurityGroupIngressResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
   <AuthorizeCacheSecurityGroupIngressResult>
      <CacheSecurityGroup>
         <EC2SecurityGroups>
            <EC2SecurityGroup>
               <Status>authorizing</Status>
               <EC2SecurityGroupName>default</EC2SecurityGroupName>
               <EC2SecurityGroupOwnerId>565419523791</EC2SecurityGroupOwnerId>
            </EC2SecurityGroup>
         </EC2SecurityGroups>
```

```
          <CacheSecurityGroupName>mygroup</CacheSecurityGroupName>
          <OwnerId>123456781234</OwnerId>
          <Description>My security group</Description>
       </CacheSecurityGroup>
    </AuthorizeCacheSecurityGroupIngress>
    <ResponseMetadata>
       <RequestId>817fa999-3647-11e0-ae57-f96cfe56749c</RequestId>
    </ResponseMetadata>
</AuthorizeCacheSecurityGroupIngressResponse>
```

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

BatchApplyUpdateAction

Apply the service update. For more information on service updates and applying them, see Applying Service Updates.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ServiceUpdateName

The unique ID of the service update

Type: String

Required: Yes

CacheClusterIds.member.N

The cache cluster IDs

Type: Array of strings

Array Members: Maximum number of 20 items.

Required: No

ReplicationGroupIds.member.N

The replication group IDs

Type: Array of strings

Array Members: Maximum number of 20 items.

Required: No

Response Elements

The following elements are returned by the service.

ProcessedUpdateActions.ProcessedUpdateAction.N

Update actions that have been processed successfully

Type: Array of ProcessedUpdateAction objects

UnprocessedUpdateActions.UnprocessedUpdateAction.N

Update actions that haven't been processed successfully

Type: Array of UnprocessedUpdateAction objects

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

ServiceUpdateNotFoundFault

The service update doesn't exist

HTTP Status Code: 404

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

BatchStopUpdateAction

Stop the service update. For more information on service updates and stopping them, see Stopping Service Updates.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ServiceUpdateName

The unique ID of the service update

Type: String

Required: Yes

CacheClusterIds.member.N

The cache cluster IDs

Type: Array of strings

Array Members: Maximum number of 20 items.

Required: No

ReplicationGroupIds.member.N

The replication group IDs

Type: Array of strings

Array Members: Maximum number of 20 items.

Required: No

Response Elements

The following elements are returned by the service.

ProcessedUpdateActions.ProcessedUpdateAction.N

Update actions that have been processed successfully

Type: Array of ProcessedUpdateAction objects

UnprocessedUpdateActions.UnprocessedUpdateAction.N

Update actions that haven't been processed successfully

Type: Array of UnprocessedUpdateAction objects

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

ServiceUpdateNotFoundFault

The service update doesn't exist

HTTP Status Code: 404

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

CompleteMigration

Complete the migration of data.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ReplicationGroupId

The ID of the replication group to which data is being migrated.

Type: String

Required: Yes

Force

Forces the migration to stop without ensuring that data is in sync. It is recommended to use this option only to abort the migration and not recommended when application wants to continue migration to ElastiCache.

Type: Boolean

Required: No

Response Elements

The following element is returned by the service.

ReplicationGroup

Contains all of the attributes of a specific Valkey or Redis OSS replication group.

Type: ReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

ReplicationGroupNotUnderMigrationFault

The designated replication group is not available for data migration.

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

CopyServerlessCacheSnapshot

Creates a copy of an existing serverless cache's snapshot. Available for Valkey, Redis OSS and Serverless Memcached only.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

SourceServerlessCacheSnapshotName

The identifier of the existing serverless cache's snapshot to be copied. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: Yes

TargetServerlessCacheSnapshotName

The identifier for the snapshot to be created. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: Yes

KmsKeyId

The identifier of the KMS key used to encrypt the target snapshot. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: No

Tags.Tag.N

A list of tags to be added to the target snapshot resource. A tag is a key-value pair. Available for Valkey, Redis OSS and Serverless Memcached only. Default: NULL

Type: Array of Tag objects

Required: No

Response Elements

The following element is returned by the service.

ServerlessCacheSnapshot

The response for the attempt to copy the serverless cache snapshot. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: ServerlessCacheSnapshot object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidServerlessCacheSnapshotStateFault

The state of the serverless cache snapshot was not received. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 400

ServerlessCacheSnapshotAlreadyExistsFault

A serverless cache snapshot with this name already exists. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 400

ServerlessCacheSnapshotNotFoundFault

This serverless cache snapshot could not be found or does not exist. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 404

ServerlessCacheSnapshotQuotaExceededFault

The number of serverless cache snapshots exceeds the customer snapshot quota. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 400

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

Examples

CopyServerlessCacheSnapshot

This example illustrates one usage of CopyServerlessCacheSnapshot.

Sample Request

```
{
    "input": {
        "SourceServerlessCacheSnapshotName": "my-serverless-cache-snapshot",
        "TargetServerlessCacheSnapshotName": "my-serverless-cache-snapshot-copy"
    },
    "output": {
        "ServerlessCacheSnapshot": {
            "ServerlessCacheSnapshotName": "my-serverless-cache-snapshot-copy",
            "ARN": "arn:aws:elasticache:us-
east-1:222222222222:serverlesscachesnapshot:my-serverless-cache-snapshot-copy",
            "SnapshotType": "manual",
            "Status": "creating",
            "ServerlessCacheConfiguration": {
                "ServerlessCacheName": "my-serverless-cache",
```

```
                "Engine": "redis",
                "MajorEngineVersion": "7"
            }
        }
    }
}
```

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

CopySnapshot

Makes a copy of an existing snapshot.

> **Note**
>
> This operation is valid for Valkey or Redis OSS only.

> **Important**
>
> Users or groups that have permissions to use the `CopySnapshot` operation can create their own Amazon S3 buckets and copy snapshots to it. To control access to your snapshots, use an IAM policy to control who has the ability to use the `CopySnapshot` operation. For more information about using IAM to control the use of ElastiCache operations, see Exporting Snapshots and Authentication & Access Control.

You could receive the following error messages.

Error Messages

- **Error Message:** The S3 bucket %s is outside of the region.

  Solution: Create an Amazon S3 bucket in the same region as your snapshot. For more information, see Step 1: Create an Amazon S3 Bucket in the ElastiCache User Guide.
- **Error Message:** The S3 bucket %s does not exist.

  Solution: Create an Amazon S3 bucket in the same region as your snapshot. For more information, see Step 1: Create an Amazon S3 Bucket in the ElastiCache User Guide.
- **Error Message:** The S3 bucket %s is not owned by the authenticated user.

  Solution: Create an Amazon S3 bucket in the same region as your snapshot. For more information, see Step 1: Create an Amazon S3 Bucket in the ElastiCache User Guide.
- **Error Message:** The authenticated user does not have sufficient permissions to perform the desired activity.

  Solution: Contact your system administrator to get the needed permissions.

- **Error Message:** The S3 bucket %s already contains an object with key %s.

  Solution: Give the `TargetSnapshotName` a new and unique value. If exporting a snapshot, you could alternatively create a new Amazon S3 bucket and use this same value for `TargetSnapshotName`.
- **Error Message:** ElastiCache has not been granted READ permissions %s on the S3 Bucket.

  Solution: Add List and Read permissions on the bucket. For more information, see Step 2: Grant ElastiCache Access to Your Amazon S3 Bucket in the ElastiCache User Guide.
- **Error Message:** ElastiCache has not been granted WRITE permissions %s on the S3 Bucket.

  Solution: Add Upload/Delete permissions on the bucket. For more information, see Step 2: Grant ElastiCache Access to Your Amazon S3 Bucket in the ElastiCache User Guide.
- **Error Message:** ElastiCache has not been granted READ_ACP permissions %s on the S3 Bucket.

  Solution: Add View Permissions on the bucket. For more information, see Step 2: Grant ElastiCache Access to Your Amazon S3 Bucket in the ElastiCache User Guide.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

SourceSnapshotName

The name of an existing snapshot from which to make a copy.

Type: String

Required: Yes

TargetSnapshotName

A name for the snapshot copy. ElastiCache does not permit overwriting a snapshot, therefore this name must be unique within its context - ElastiCache or an Amazon S3 bucket if exporting.

Type: String

Required: Yes

KmsKeyId

The ID of the KMS key used to encrypt the target snapshot.

Type: String

Required: No

Tags.Tag.N

A list of tags to be added to this resource. A tag is a key-value pair. A tag key must be accompanied by a tag value, although null is accepted.

Type: Array of Tag objects

Required: No

TargetBucket

The Amazon S3 bucket to which the snapshot is exported. This parameter is used only when exporting a snapshot for external access.

When using this parameter to export a snapshot, be sure Amazon ElastiCache has the needed permissions to this S3 bucket. For more information, see Step 2: Grant ElastiCache Access to Your Amazon S3 Bucket in the *Amazon ElastiCache User Guide*.

For more information, see Exporting a Snapshot in the *Amazon ElastiCache User Guide*.

Type: String

Required: No

Response Elements

The following element is returned by the service.

Snapshot

Represents a copy of an entire Valkey or Redis OSS cluster as of the time when the snapshot was taken.

Type: Snapshot object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidSnapshotState

The current state of the snapshot does not allow the requested operation to occur.

HTTP Status Code: 400

SnapshotAlreadyExistsFault

You already have a snapshot with the given name.

HTTP Status Code: 400

SnapshotNotFoundFault

The requested snapshot name does not refer to an existing snapshot.

HTTP Status Code: 404

SnapshotQuotaExceededFault

The request cannot be processed because it would exceed the maximum number of snapshots.

HTTP Status Code: 400

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

Examples

Snapshot copy

The following example makes a copy of the snapshot `automatic.my-redis-primary-2016-04-27-03-15` named `my-snapshot-copy`.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
    ?Action=CopySnapshot
    &SourceSnapshotName=automatic.my-redis-primary-2016-04-27-03-15
    &TargetSnapshotName=my-snapshot-copy
    &SignatureVersion=4
    &SignatureMethod=HmacSHA256
    &Timestamp=20141201T220302Z
    &Version=2015-02-02
    &X-Amz-Algorithm=AWS4-HMAC-SHA256
    &X-Amz-Date=20141201T220302Z
    &X-Amz-SignedHeaders=Host
    &X-Amz-Expires=20141201T220302Z
    &X-Amz-Credential=<credential>
    &X-Amz-Signature=<signature>
```

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

CreateCacheCluster

Creates a cluster. All nodes in the cluster run the same protocol-compliant cache engine software, either Memcached, Valkey or Redis OSS.

This operation is not supported for Valkey or Redis OSS (cluster mode enabled) clusters.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheClusterId

The node group (shard) identifier. This parameter is stored as a lowercase string.

Constraints:

- A name must contain from 1 to 50 alphanumeric characters or hyphens.
- The first character must be a letter.
- A name cannot end with a hyphen or contain two consecutive hyphens.

Type: String

Required: Yes

AuthToken

Reserved parameter. The password used to access a password protected server.

Password constraints:

- Must be only printable ASCII characters.
- Must be at least 16 characters and no more than 128 characters in length.
- The only permitted printable special characters are !, &, #, $, ^, <, >, and -. Other printable special characters cannot be used in the AUTH token.

For more information, see AUTH password at http://redis.io/commands/AUTH.

Type: String

Required: No

AutoMinorVersionUpgrade

If you are running Valkey 7.2 and above or Redis OSS engine version 6.0 and above, set this parameter to yes to opt-in to the next auto minor version upgrade campaign. This parameter is disabled for previous versions.

Type: Boolean

Required: No

AZMode

Specifies whether the nodes in this Memcached cluster are created in a single Availability Zone or created across multiple Availability Zones in the cluster's region.

This parameter is only supported for Memcached clusters.

If the `AZMode` and `PreferredAvailabilityZones` are not specified, ElastiCache assumes `single-az` mode.

Type: String

Valid Values: `single-az | cross-az`

Required: No

CacheNodeType

The compute and memory capacity of the nodes in the node group (shard).

The following node types are supported by ElastiCache. Generally speaking, the current generation types provide more memory and computational power at lower cost when compared to their equivalent previous generation counterparts.

- General purpose:
  - Current generation:

    M7g node types: `cache.m7g.large`, `cache.m7g.xlarge`, `cache.m7g.2xlarge`, `cache.m7g.4xlarge`, `cache.m7g.8xlarge`, `cache.m7g.12xlarge`, `cache.m7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

M6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.m6g.large`, `cache.m6g.xlarge`, `cache.m6g.2xlarge`, `cache.m6g.4xlarge`, `cache.m6g.8xlarge`, `cache.m6g.12xlarge`, `cache.m6g.16xlarge`

M5 node types: `cache.m5.large`, `cache.m5.xlarge`, `cache.m5.2xlarge`, `cache.m5.4xlarge`, `cache.m5.12xlarge`, `cache.m5.24xlarge`

M4 node types: `cache.m4.large`, `cache.m4.xlarge`, `cache.m4.2xlarge`, `cache.m4.4xlarge`, `cache.m4.10xlarge`

T4g node types (available only for Redis OSS engine version 5.0.6 onward and Memcached engine version 1.5.16 onward): `cache.t4g.micro`, `cache.t4g.small`, `cache.t4g.medium`

T3 node types: `cache.t3.micro`, `cache.t3.small`, `cache.t3.medium`

T2 node types: `cache.t2.micro`, `cache.t2.small`, `cache.t2.medium`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  T1 node types: `cache.t1.micro`

  M1 node types: `cache.m1.small`, `cache.m1.medium`, `cache.m1.large`, `cache.m1.xlarge`

  M3 node types: `cache.m3.medium`, `cache.m3.large`, `cache.m3.xlarge`, `cache.m3.2xlarge`

- Compute optimized:
  - Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

    C1 node types: `cache.c1.xlarge`

- Memory optimized:
  - Current generation:

R7g node types: `cache.r7g.large`, `cache.r7g.xlarge`, `cache.r7g.2xlarge`, `cache.r7g.4xlarge`, `cache.r7g.8xlarge`, `cache.r7g.12xlarge`, `cache.r7g.16xlarge`

> **Note**
>
> For region availability, see Supported Node Types

R6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.r6g.large`, `cache.r6g.xlarge`, `cache.r6g.2xlarge`, `cache.r6g.4xlarge`, `cache.r6g.8xlarge`, `cache.r6g.12xlarge`, `cache.r6g.16xlarge`

R5 node types: `cache.r5.large`, `cache.r5.xlarge`, `cache.r5.2xlarge`, `cache.r5.4xlarge`, `cache.r5.12xlarge`, `cache.r5.24xlarge`

R4 node types: `cache.r4.large`, `cache.r4.xlarge`, `cache.r4.2xlarge`, `cache.r4.4xlarge`, `cache.r4.8xlarge`, `cache.r4.16xlarge`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  M2 node types: `cache.m2.xlarge`, `cache.m2.2xlarge`, `cache.m2.4xlarge`

  R3 node types: `cache.r3.large`, `cache.r3.xlarge`, `cache.r3.2xlarge`, `cache.r3.4xlarge`, `cache.r3.8xlarge`

Additional node type info

- All current generation instance types are created in Amazon VPC by default.
- Valkey or Redis OSS append-only files (AOF) are not supported for T1 or T2 instances.
- Valkey or Redis OSS Multi-AZ with automatic failover is not supported on T1 instances.
- The configuration variables `appendonly` and `appendfsync` are not supported on Valkey, or on Redis OSS version 2.8.22 and later.

Type: String

Required: No

CacheParameterGroupName

The name of the parameter group to associate with this cluster. If this argument is omitted, the default parameter group for the specified engine is used. You cannot use any parameter group which has `cluster-enabled='yes'` when creating a cluster.

Type: String

Required: No

CacheSecurityGroupNames.CacheSecurityGroupName.N

A list of security group names to associate with this cluster.

Use this parameter only when you are creating a cluster outside of an Amazon Virtual Private Cloud (Amazon VPC).

Type: Array of strings

Required: No

CacheSubnetGroupName

The name of the subnet group to be used for the cluster.

Use this parameter only when you are creating a cluster in an Amazon Virtual Private Cloud (Amazon VPC).

> **Important**
>
> If you're going to launch your cluster in an Amazon VPC, you need to create a subnet group before you start creating a cluster. For more information, see Subnets and Subnet Groups.

Type: String

Required: No

Engine

The name of the cache engine to be used for this cluster.

Valid values for this parameter are: `memcached` | `redis`

Type: String

Required: No

EngineVersion

The version number of the cache engine to be used for this cluster. To view the supported cache engine versions, use the DescribeCacheEngineVersions operation.

Important: You can upgrade to a newer engine version (see Selecting a Cache Engine and Version), but you cannot downgrade to an earlier engine version. If you want to use an earlier engine version, you must delete the existing cluster or replication group and create it anew with the earlier engine version.

Type: String

Required: No

IpDiscovery

The network type you choose when modifying a cluster, either `ipv4` | `ipv6`. IPv6 is supported for workloads using Valkey 7.2 and above, Redis OSS engine version 6.2 to 7.1 and Memcached engine version 1.6.6 and above on all instances built on the Nitro system.

Type: String

Valid Values: `ipv4 | ipv6`

Required: No

LogDeliveryConfigurations.LogDeliveryConfigurationRequest.N

Specifies the destination, format and type of the logs.

Type: Array of LogDeliveryConfigurationRequest objects

Required: No

NetworkType

Must be either `ipv4` | `ipv6` | `dual_stack`. IPv6 is supported for workloads using Valkey 7.2 and above, Redis OSS engine version 6.2 to 7.1 and Memcached engine version 1.6.6 and above on all instances built on the Nitro system.

Type: String

Valid Values: `ipv4 | ipv6 | dual_stack`

Required: No

NotificationTopicArn

The Amazon Resource Name (ARN) of the Amazon Simple Notification Service (SNS) topic to which notifications are sent.

> **Note**
>
> The Amazon SNS topic owner must be the same as the cluster owner.

Type: String

Required: No

NumCacheNodes

The initial number of cache nodes that the cluster has.

For clusters running Valkey or Redis OSS, this value must be 1. For clusters running Memcached, this value must be between 1 and 40.

If you need more than 40 nodes for your Memcached cluster, please fill out the ElastiCache Limit Increase Request form at http://aws.amazon.com/contact-us/elasticache-node-limit-request/.

Type: Integer

Required: No

OutpostMode

Specifies whether the nodes in the cluster are created in a single outpost or across multiple outposts.

Type: String

Valid Values: `single-outpost | cross-outpost`

Required: No

Port

The port number on which each of the cache nodes accepts connections.

Type: Integer

Required: No

PreferredAvailabilityZone

The EC2 Availability Zone in which the cluster is created.

All nodes belonging to this cluster are placed in the preferred Availability Zone. If you want to create your nodes across multiple Availability Zones, use `PreferredAvailabilityZones`.

Default: System chosen Availability Zone.

Type: String

Required: No

PreferredAvailabilityZones.PreferredAvailabilityZone.N

A list of the Availability Zones in which cache nodes are created. The order of the zones in the list is not important.

This option is only supported on Memcached.

> **Note**
>
> If you are creating your cluster in an Amazon VPC (recommended) you can only locate nodes in Availability Zones that are associated with the subnets in the selected subnet group.
> The number of Availability Zones listed must equal the value of `NumCacheNodes`.

If you want all the nodes in the same Availability Zone, use `PreferredAvailabilityZone` instead, or repeat the Availability Zone multiple times in the list.

Default: System chosen Availability Zones.

Type: Array of strings

Required: No

PreferredMaintenanceWindow

Specifies the weekly time range during which maintenance on the cluster is performed. It is specified as a range in the format ddd:hh24:mi-ddd:hh24:mi (24H Clock UTC). The minimum maintenance window is a 60 minute period.

Type: String

Required: No

PreferredOutpostArn

The outpost ARN in which the cache cluster is created.

Type: String

Required: No

PreferredOutpostArns.PreferredOutpostArn.N

The outpost ARNs in which the cache cluster is created.

Type: Array of strings

Required: No

ReplicationGroupId

The ID of the replication group to which this cluster should belong. If this parameter is specified, the cluster is added to the specified replication group as a read replica; otherwise, the cluster is a standalone primary that is not part of any replication group.

If the specified replication group is Multi-AZ enabled and the Availability Zone is not specified, the cluster is created in Availability Zones that provide the best spread of read replicas across Availability Zones.

> **Note**
>
> This parameter is only valid if the `Engine` parameter is `redis`.

Type: String

Required: No

SecurityGroupIds.SecurityGroupId.N

One or more VPC security groups associated with the cluster.

Use this parameter only when you are creating a cluster in an Amazon Virtual Private Cloud (Amazon VPC).

Type: Array of strings

Required: No

SnapshotArns.SnapshotArn.N

A single-element string list containing an Amazon Resource Name (ARN) that uniquely identifies a Valkey or Redis OSS RDB snapshot file stored in Amazon S3. The snapshot file is used to populate the node group (shard). The Amazon S3 object name in the ARN cannot contain any commas.

> **Note**
>
> This parameter is only valid if the `Engine` parameter is `redis`.

Example of an Amazon S3 ARN: `arn:aws:s3:::my_bucket/snapshot1.rdb`

Type: Array of strings

Required: No

SnapshotName

The name of a Valkey or Redis OSS snapshot from which to restore data into the new node group (shard). The snapshot status changes to `restoring` while the new node group (shard) is being created.

> **Note**
>
> This parameter is only valid if the `Engine` parameter is `redis`.

Type: String

Required: No

SnapshotRetentionLimit

The number of days for which ElastiCache retains automatic snapshots before deleting them. For example, if you set `SnapshotRetentionLimit` to 5, a snapshot taken today is retained for 5 days before being deleted.

> **Note**
>
> This parameter is only valid if the `Engine` parameter is `redis`.

Default: 0 (i.e., automatic backups are disabled for this cache cluster).

Type: Integer

Required: No

SnapshotWindow

The daily time range (in UTC) during which ElastiCache begins taking a daily snapshot of your node group (shard).

Example: `05:00-09:00`

If you do not specify this parameter, ElastiCache automatically chooses an appropriate time range.

> **Note**
>
> This parameter is only valid if the `Engine` parameter is `redis`.

Type: String

Required: No

Tags.Tag.N

A list of tags to be added to this resource.

Type: Array of Tag objects

Required: No

TransitEncryptionEnabled

A flag that enables in-transit encryption when set to true.

Type: Boolean

Required: No

Response Elements

The following element is returned by the service.

CacheCluster

Contains all of the attributes of a specific cluster.

Type: CacheCluster object

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterAlreadyExists

You already have a cluster with the given identifier.

HTTP Status Code: 400

CacheParameterGroupNotFound

The requested cache parameter group name does not refer to an existing cache parameter group.

HTTP Status Code: 404

CacheSecurityGroupNotFound

The requested cache security group name does not refer to an existing cache security group.

HTTP Status Code: 404

CacheSubnetGroupNotFoundFault

The requested cache subnet group name does not refer to an existing cache subnet group.

HTTP Status Code: 400

ClusterQuotaForCustomerExceeded

The request cannot be processed because it would exceed the allowed number of clusters per customer.

HTTP Status Code: 400

InsufficientCacheClusterCapacity

The requested cache node type is not available in the specified Availability Zone. For more information, see InsufficientCacheClusterCapacity in the ElastiCache User Guide.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

InvalidVPCNetworkStateFault

The VPC network is in an invalid state.

HTTP Status Code: 400

NodeQuotaForClusterExceeded

The request cannot be processed because it would exceed the allowed number of cache nodes in a single cluster.

HTTP Status Code: 400

NodeQuotaForCustomerExceeded

The request cannot be processed because it would exceed the allowed number of cache nodes per customer.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

Examples

CreateCacheCluster

This example illustrates one usage of CreateCacheCluster.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
 ?Action=CreateCacheCluster
 &CacheClusterId=myMemcachedCluster
 &CacheNodeType=cache.m1.small
 &CacheSecurityGroupNames.CacheSecurityGroupName.1=default
 &Engine=memcached
 &NumCacheNodes=3
 &PreferredAvailabilityZones.PreferredAvailabilityZone.1=us-west-2a
 &PreferredAvailabilityZones.PreferredAvailabilityZone.2=us-west-2b
 &PreferredAvailabilityZones.PreferredAvailabilityZone.3=us-west-2c
 &SignatureMethod=HmacSHA256
 &SignatureVersion=4
 &Version=2015-02-02
 &X-Amz-Algorithm=AWS4-HMAC-SHA256
 &X-Amz-Credential=[your-access-key-id]/20150202/us-west-2/elasticache/AWS4_request
```

```
&X-Amz-Date=20150202T170651Z
&X-Amz-SignedHeaders=content-type;host;user-agent;x-amz-content-sha256;x-amz-date
&X-Amz-Signature=[signature-value]
```

Sample Response

```
<CreateCacheClusterResponse xmlns="http://elasticache.amazonaws.com/doc/2015-02-02/">
  <CreateCacheClusterResult>
    <CacheCluster>
      <CacheClusterId>myMemcachedClustger</CacheClusterId>
      <CacheClusterStatus>creating</CacheClusterStatus>
      <CacheParameterGroup>
        <CacheParameterGroupName>default.memcached1.4</CacheParameterGroupName>
        <ParameterApplyStatus>in-sync</ParameterApplyStatus>
        <CacheNodeIdsToReboot/>
      </CacheParameterGroup>
      <CacheNodeType>cache.m1.small</CacheNodeType>
      <Engine>memcached</Engine>
      <PendingModifiedValues/>
      <EngineVersion>1.4.14</EngineVersion>
      <AutoMinorVersionUpgrade>true</AutoMinorVersionUpgrade>
      <PreferredMaintenanceWindow>sat:09:00-sat:10:00</PreferredMaintenanceWindow>
      <ClientDownloadLandingPage>https://console.aws.amazon.com/elasticache/
home#client-download:</ClientDownloadLandingPage>
      <CacheSecurityGroups>
        <CacheSecurityGroup>
          <CacheSecurityGroupName>default</CacheSecurityGroupName>
          <Status>active</Status>
        </CacheSecurityGroup>
      </CacheSecurityGroups>
      <NumCacheNodes>3</NumCacheNodes>
    </CacheCluster>
  </CreateCacheClusterResult>
  <ResponseMetadata>
    <RequestId>69134921-10f9-11e4-81bb-d76bad68b8fd</RequestId>
  </ResponseMetadata>
</CreateCacheClusterResponse>
```

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

CreateCacheParameterGroup

Creates a new Amazon ElastiCache cache parameter group. An ElastiCache cache parameter group is a collection of parameters and their values that are applied to all of the nodes in any cluster or replication group using the CacheParameterGroup.

A newly created CacheParameterGroup is an exact duplicate of the default parameter group for the CacheParameterGroupFamily. To customize the newly created CacheParameterGroup you can change the values of specific parameters. For more information, see:

- ModifyCacheParameterGroup in the ElastiCache API Reference.
- Parameters and Parameter Groups in the ElastiCache User Guide.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheParameterGroupFamily

The name of the cache parameter group family that the cache parameter group can be used with.

Valid values are: `valkey8` | `valkey7` | `memcached1.4` | `memcached1.5` | `memcached1.6` | `redis2.6` | `redis2.8` | `redis3.2` | `redis4.0` | `redis5.0` | `redis6.x` | `redis7`

Type: String

Required: Yes

CacheParameterGroupName

A user-specified name for the cache parameter group.

Type: String

Required: Yes

Description

A user-specified description for the cache parameter group.

Type: String

Required: Yes

Tags.Tag.N

A list of tags to be added to this resource. A tag is a key-value pair. A tag key must be accompanied by a tag value, although null is accepted.

Type: Array of Tag objects

Required: No

Response Elements

The following element is returned by the service.

CacheParameterGroup

Represents the output of a `CreateCacheParameterGroup` operation.

Type: CacheParameterGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheParameterGroupAlreadyExists

A cache parameter group with the requested name already exists.

HTTP Status Code: 400

CacheParameterGroupQuotaExceeded

The request cannot be processed because it would exceed the maximum number of cache security groups.

HTTP Status Code: 400

InvalidCacheParameterGroupState

The current state of the cache parameter group does not allow the requested operation to occur.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

Examples

CreateCacheParameterGroup

This example illustrates one usage of CreateCacheParameterGroup.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=CreateCacheParameterGroup
   &CacheParameterGroupFamily=memcached1.4
   &CacheParameterGroupName=mycacheparametergroup1
   &Description=My%20custom%20Redis%20cache%20parameter%20group
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &Version=2015-02-02
   &X-Amz-Credential=<credential>
```

Sample Response

```
<CreateCacheParameterGroupResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
  <CreateCacheParameterGroupResult>
    <CacheParameterGroup>
      <CacheParameterGroupName>mycacheparametergroup1</CacheParameterGroupName>
      <CacheParameterGroupFamily>memcached1.4</CacheParameterGroupFamily>
      <Description>My first cache parameter group</Description>
    </CacheParameterGroup>
  </CreateCacheParameterGroupResult>
  <ResponseMetadata>
    <RequestId>05699541-b7f9-11e0-9326-b7275b9d4a6c</RequestId>
  </ResponseMetadata>
</CreateCacheParameterGroupResponse>
```

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

CreateCacheSecurityGroup

Creates a new cache security group. Use a cache security group to control access to one or more clusters.

Cache security groups are only used when you are creating a cluster outside of an Amazon Virtual Private Cloud (Amazon VPC). If you are creating a cluster inside of a VPC, use a cache subnet group instead. For more information, see CreateCacheSubnetGroup.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheSecurityGroupName

A name for the cache security group. This value is stored as a lowercase string.

Constraints: Must contain no more than 255 alphanumeric characters. Cannot be the word "Default".

Example: `mysecuritygroup`

Type: String

Required: Yes

Description

A description for the cache security group.

Type: String

Required: Yes

Tags.Tag.N

A list of tags to be added to this resource. A tag is a key-value pair. A tag key must be accompanied by a tag value, although null is accepted.

Type: Array of Tag objects

Required: No

Response Elements

The following element is returned by the service.

CacheSecurityGroup

Represents the output of one of the following operations:

- `AuthorizeCacheSecurityGroupIngress`
- `CreateCacheSecurityGroup`
- `RevokeCacheSecurityGroupIngress`

Type: CacheSecurityGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheSecurityGroupAlreadyExists

A cache security group with the specified name already exists.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

QuotaExceeded.CacheSecurityGroup

The request cannot be processed because it would exceed the allowed number of cache security groups.

HTTP Status Code: 400

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

Examples

CreateCacheSecurityGroup

This example illustrates one usage of CreateCacheSecurityGroup.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=CreateCacheSecurityGroup
   &CacheSecurityGroupName=mycachesecuritygroup
   &Description=My%20cache%20security%20group
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<CreateCacheSecurityGroupResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
   <CreateCacheSecurityGroupResult>
      <CacheSecurityGroup>
         <EC2SecurityGroups/>
         <CacheSecurityGroupName>mycachesecuritygroup</CacheSecurityGroupName>
         <OwnerId>123456789012</OwnerId>
         <Description>My cache security group</Description>
      </CacheSecurityGroup>
   </CreateCacheSecurityGroupResult>
   <ResponseMetadata>
      <RequestId>2b1c8035-b7fa-11e0-9326-b7275b9d4a6c</RequestId>
   </ResponseMetadata>
</CreateCacheSecurityGroupResponse>
```

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

CreateCacheSubnetGroup

Creates a new cache subnet group.

Use this parameter only when you are creating a cluster in an Amazon Virtual Private Cloud (Amazon VPC).

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheSubnetGroupDescription

A description for the cache subnet group.

Type: String

Required: Yes

CacheSubnetGroupName

A name for the cache subnet group. This value is stored as a lowercase string.

Constraints: Must contain no more than 255 alphanumeric characters or hyphens.

Example: `mysubnetgroup`

Type: String

Required: Yes

SubnetIds.SubnetIdentifier.N

A list of VPC subnet IDs for the cache subnet group.

Type: Array of strings

Required: Yes

Tags.Tag.N

A list of tags to be added to this resource. A tag is a key-value pair. A tag key must be accompanied by a tag value, although null is accepted.

Type: Array of Tag objects

Required: No

Response Elements

The following element is returned by the service.

CacheSubnetGroup

Represents the output of one of the following operations:

- `CreateCacheSubnetGroup`
- `ModifyCacheSubnetGroup`

Type: CacheSubnetGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheSubnetGroupAlreadyExists

The requested cache subnet group name is already in use by an existing cache subnet group.

HTTP Status Code: 400

CacheSubnetGroupQuotaExceeded

The request cannot be processed because it would exceed the allowed number of cache subnet groups.

HTTP Status Code: 400

CacheSubnetQuotaExceededFault

The request cannot be processed because it would exceed the allowed number of subnets in a cache subnet group.

HTTP Status Code: 400

InvalidSubnet

An invalid subnet identifier was specified.

HTTP Status Code: 400

SubnetNotAllowedFault

At least one subnet ID does not match the other subnet IDs. This mismatch typically occurs when a user sets one subnet ID to a regional Availability Zone and a different one to an outpost. Or when a user sets the subnet ID to an Outpost when not subscribed on this service.

HTTP Status Code: 400

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

Examples

CreateCacheSubnetGroup

This example illustrates one usage of CreateCacheSubnetGroup.

Sample Request

```
https://elasticache.amazonaws.com/
   ?Action=CreateCacheSubnetGroup
   &CacheSubnetGroupName=myCachesubnetgroup
   &CacheSubnetGroupDescription=My%20new%20CacheSubnetGroup
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<CreateCacheSubnetGroupResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
   <CreateCacheSubnetGroupResult>
      <CacheSubnetGroup>
         <VpcId>990524496922</VpcId>
```

```
            <CacheSubnetGroupDescription>My new CacheSubnetGroup</
CacheSubnetGroupDescription>
            <CacheSubnetGroupName>myCachesubnetgroup</CacheSubnetGroupName>
            <Subnets>
               <Subnet>
                  <SubnetStatus>Active</SubnetStatus>
                  <SubnetIdentifier>subnet-7c5b4115</SubnetIdentifier>
                  <SubnetAvailabilityZone>
                     <Name>us-west-2c</Name>
                  </SubnetAvailabilityZone>
               </Subnet>
               <Subnet>
                  <SubnetStatus>Active</SubnetStatus>
                  <SubnetIdentifier>subnet-7b5b4112</SubnetIdentifier>
                  <SubnetAvailabilityZone>
                     <Name>us-west-2b</Name>
                  </SubnetAvailabilityZone>
               </Subnet>
               <Subnet>
                  <SubnetStatus>Active</SubnetStatus>
                  <SubnetIdentifier>subnet-3ea6bd57</SubnetIdentifier>
                  <SubnetAvailabilityZone>
                     <Name>us-west-2c</Name>
                  </SubnetAvailabilityZone>
               </Subnet>
            </Subnets>
         </CacheSubnetGroup>
      </CreateCacheSubnetGroupResult>
      <ResponseMetadata>
         <RequestId>ed662948-a57b-11df-9e38-7ffab86c801f</RequestId>
      </ResponseMetadata>
   </CreateCacheSubnetGroupResponse>
```

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

CreateGlobalReplicationGroup

Global Datastore offers fully managed, fast, reliable and secure cross-region replication. Using Global Datastore with Valkey or Redis OSS, you can create cross-region read replica clusters for ElastiCache to enable low-latency reads and disaster recovery across regions. For more information, see Replication Across Regions Using Global Datastore.

- The **GlobalReplicationGroupIdSuffix** is the name of the Global datastore.
- The **PrimaryReplicationGroupId** represents the name of the primary cluster that accepts writes and will replicate updates to the secondary cluster.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

GlobalReplicationGroupIdSuffix

The suffix name of a Global datastore. Amazon ElastiCache automatically applies a prefix to the Global datastore ID when it is created. Each Amazon Region has its own prefix. For instance, a Global datastore ID created in the US-West-1 region will begin with "dsdfu" along with the suffix name you provide. The suffix, combined with the auto-generated prefix, guarantees uniqueness of the Global datastore name across multiple regions.

For a full list of Amazon Regions and their respective Global datastore iD prefixes, see Using the Amazon CLI with Global datastores .

Type: String

Required: Yes

PrimaryReplicationGroupId

The name of the primary cluster that accepts writes and will replicate updates to the secondary cluster.

Type: String

Required: Yes

GlobalReplicationGroupDescription

Provides details of the Global datastore

Type: String

Required: No

Response Elements

The following element is returned by the service.

GlobalReplicationGroup

Consists of a primary cluster that accepts writes and an associated secondary cluster that resides in a different Amazon region. The secondary cluster accepts only reads. The primary cluster automatically replicates updates to the secondary cluster.

- The **GlobalReplicationGroupIdSuffix** represents the name of the Global datastore, which is what you use to associate a secondary cluster.

Type: GlobalReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

GlobalReplicationGroupAlreadyExistsFault

The Global datastore name already exists.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

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

CreateReplicationGroup

Creates a Valkey or Redis OSS (cluster mode disabled) or a Valkey or Redis OSS (cluster mode enabled) replication group.

This API can be used to create a standalone regional replication group or a secondary replication group associated with a Global datastore.

A Valkey or Redis OSS (cluster mode disabled) replication group is a collection of nodes, where one of the nodes is a read/write primary and the others are read-only replicas. Writes to the primary are asynchronously propagated to the replicas.

A Valkey or Redis OSS cluster-mode enabled cluster is comprised of from 1 to 90 shards (API/CLI: node groups). Each shard has a primary node and up to 5 read-only replica nodes. The configuration can range from 90 shards and 0 replicas to 15 shards and 5 replicas, which is the maximum number or replicas allowed.

The node or shard limit can be increased to a maximum of 500 per cluster if the Valkey or Redis OSS engine version is 5.0.6 or higher. For example, you can choose to configure a 500 node cluster that ranges between 83 shards (one primary and 5 replicas per shard) and 500 shards (single primary and no replicas). Make sure there are enough available IP addresses to accommodate the increase. Common pitfalls include the subnets in the subnet group have too small a CIDR range or the subnets are shared and heavily used by other clusters. For more information, see Creating a Subnet Group. For versions below 5.0.6, the limit is 250 per cluster.

To request a limit increase, see Amazon Service Limits and choose the limit type **Nodes per cluster per instance type**.

When a Valkey or Redis OSS (cluster mode disabled) replication group has been successfully created, you can add one or more read replicas to it, up to a total of 5 read replicas. If you need to increase or decrease the number of node groups (console: shards), you can use scaling. For more information, see Scaling self-designed clusters in the *ElastiCache User Guide*.

> **Note**
>
> This operation is valid for Valkey and Redis OSS only.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ReplicationGroupDescription

A user-created description for the replication group.

Type: String

Required: Yes

ReplicationGroupId

The replication group identifier. This parameter is stored as a lowercase string.

Constraints:

- A name must contain from 1 to 40 alphanumeric characters or hyphens.
- The first character must be a letter.
- A name cannot end with a hyphen or contain two consecutive hyphens.

Type: String

Required: Yes

AtRestEncryptionEnabled

A flag that enables encryption at rest when set to `true`.

You cannot modify the value of `AtRestEncryptionEnabled` after the replication group is created. To enable encryption at rest on a replication group you must set `AtRestEncryptionEnabled` to `true` when you create the replication group.

Required: Only available when creating a replication group in an Amazon VPC using Valkey 7.2 and later, Redis OSS version `3.2.6`, or Redis OSS `4.x` and later.

Default: `true` when using Valkey, `false` when using Redis OSS

Type: Boolean

Required: No

AuthToken

Reserved parameter. The password used to access a password protected server.

`AuthToken` can be specified only on replication groups where `TransitEncryptionEnabled` is `true`.

> **Important**
>
> For HIPAA compliance, you must specify `TransitEncryptionEnabled` as `true`, an `AuthToken`, and a `CacheSubnetGroup`.

Password constraints:

- Must be only printable ASCII characters.
- Must be at least 16 characters and no more than 128 characters in length.
- The only permitted printable special characters are !, &, #, $, ^, <, >, and -. Other printable special characters cannot be used in the AUTH token.

For more information, see AUTH password at http://redis.io/commands/AUTH.

Type: String

Required: No

AutomaticFailoverEnabled

Specifies whether a read-only replica is automatically promoted to read/write primary if the existing primary fails.

`AutomaticFailoverEnabled` must be enabled for Valkey or Redis OSS (cluster mode enabled) replication groups.

Default: false

Type: Boolean

Required: No

AutoMinorVersionUpgrade

If you are running Valkey 7.2 and above or Redis OSS engine version 6.0 and above, set this parameter to yes to opt-in to the next auto minor version upgrade campaign. This parameter is disabled for previous versions.

Type: Boolean

Required: No

CacheNodeType

The compute and memory capacity of the nodes in the node group (shard).

The following node types are supported by ElastiCache. Generally speaking, the current generation types provide more memory and computational power at lower cost when compared to their equivalent previous generation counterparts.

- General purpose:
  - Current generation:

    M7g node types: `cache.m7g.large`, `cache.m7g.xlarge`, `cache.m7g.2xlarge`, `cache.m7g.4xlarge`, `cache.m7g.8xlarge`, `cache.m7g.12xlarge`, `cache.m7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    M6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.m6g.large`, `cache.m6g.xlarge`, `cache.m6g.2xlarge`, `cache.m6g.4xlarge`, `cache.m6g.8xlarge`, `cache.m6g.12xlarge`, `cache.m6g.16xlarge`

    M5 node types: `cache.m5.large`, `cache.m5.xlarge`, `cache.m5.2xlarge`, `cache.m5.4xlarge`, `cache.m5.12xlarge`, `cache.m5.24xlarge`

    M4 node types: `cache.m4.large`, `cache.m4.xlarge`, `cache.m4.2xlarge`, `cache.m4.4xlarge`, `cache.m4.10xlarge`

    T4g node types (available only for Redis OSS engine version 5.0.6 onward and Memcached engine version 1.5.16 onward): `cache.t4g.micro`, `cache.t4g.small`, `cache.t4g.medium`

    T3 node types: `cache.t3.micro`, `cache.t3.small`, `cache.t3.medium`

T2 node types: `cache.t2.micro`, `cache.t2.small`, `cache.t2.medium`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  T1 node types: `cache.t1.micro`

  M1 node types: `cache.m1.small`, `cache.m1.medium`, `cache.m1.large`, `cache.m1.xlarge`

  M3 node types: `cache.m3.medium`, `cache.m3.large`, `cache.m3.xlarge`, `cache.m3.2xlarge`

- Compute optimized:
  - Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

    C1 node types: `cache.c1.xlarge`

- Memory optimized:
  - Current generation:

    R7g node types: `cache.r7g.large`, `cache.r7g.xlarge`, `cache.r7g.2xlarge`, `cache.r7g.4xlarge`, `cache.r7g.8xlarge`, `cache.r7g.12xlarge`, `cache.r7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    R6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.r6g.large`, `cache.r6g.xlarge`, `cache.r6g.2xlarge`, `cache.r6g.4xlarge`, `cache.r6g.8xlarge`, `cache.r6g.12xlarge`, `cache.r6g.16xlarge`

    R5 node types: `cache.r5.large`, `cache.r5.xlarge`, `cache.r5.2xlarge`, `cache.r5.4xlarge`, `cache.r5.12xlarge`, `cache.r5.24xlarge`

R4 node types: `cache.r4.large`, `cache.r4.xlarge`, `cache.r4.2xlarge`, `cache.r4.4xlarge`, `cache.r4.8xlarge`, `cache.r4.16xlarge`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  M2 node types: `cache.m2.xlarge`, `cache.m2.2xlarge`, `cache.m2.4xlarge`

  R3 node types: `cache.r3.large`, `cache.r3.xlarge`, `cache.r3.2xlarge`, `cache.r3.4xlarge`, `cache.r3.8xlarge`

Additional node type info

- All current generation instance types are created in Amazon VPC by default.
- Valkey or Redis OSS append-only files (AOF) are not supported for T1 or T2 instances.
- Valkey or Redis OSS Multi-AZ with automatic failover is not supported on T1 instances.
- The configuration variables `appendonly` and `appendfsync` are not supported on Valkey, or on Redis OSS version 2.8.22 and later.

Type: String

Required: No

CacheParameterGroupName

The name of the parameter group to associate with this replication group. If this argument is omitted, the default cache parameter group for the specified engine is used.

If you are running Valkey or Redis OSS version 3.2.4 or later, only one node group (shard), and want to use a default parameter group, we recommend that you specify the parameter group by name.

- To create a Valkey or Redis OSS (cluster mode disabled) replication group, use `CacheParameterGroupName=default.redis3.2`.
- To create a Valkey or Redis OSS (cluster mode enabled) replication group, use `CacheParameterGroupName=default.redis3.2.cluster.on`.

Type: String

Required: No

CacheSecurityGroupNames.CacheSecurityGroupName.N

A list of cache security group names to associate with this replication group.

Type: Array of strings

Required: No

CacheSubnetGroupName

The name of the cache subnet group to be used for the replication group.

> **Important**
>
> If you're going to launch your cluster in an Amazon VPC, you need to create a subnet group before you start creating a cluster. For more information, see Subnets and Subnet Groups.

Type: String

Required: No

ClusterMode

Enabled or Disabled. To modify cluster mode from Disabled to Enabled, you must first set the cluster mode to Compatible. Compatible mode allows your Valkey or Redis OSS clients to connect using both cluster mode enabled and cluster mode disabled. After you migrate all Valkey or Redis OSS clients to use cluster mode enabled, you can then complete cluster mode configuration and set the cluster mode to Enabled.

Type: String

Valid Values: `enabled | disabled | compatible`

Required: No

DataTieringEnabled

Enables data tiering. Data tiering is only supported for replication groups using the r6gd node type. This parameter must be set to true when using r6gd nodes. For more information, see Data tiering.

Type: Boolean

Required: No

Engine

The name of the cache engine to be used for the clusters in this replication group. The value must be set to `valkey` or `redis`.

Type: String

Required: No

EngineVersion

The version number of the cache engine to be used for the clusters in this replication group. To view the supported cache engine versions, use the `DescribeCacheEngineVersions` operation.

Important: You can upgrade to a newer engine version (see Selecting a Cache Engine and Version) in the *ElastiCache User Guide*, but you cannot downgrade to an earlier engine version. If you want to use an earlier engine version, you must delete the existing cluster or replication group and create it anew with the earlier engine version.

Type: String

Required: No

GlobalReplicationGroupId

The name of the Global datastore

Type: String

Required: No

IpDiscovery

The network type you choose when creating a replication group, either `ipv4 | ipv6`. IPv6 is supported for workloads using Valkey 7.2 and above, Redis OSS engine version 6.2 to 7.1 or Memcached engine version 1.6.6 and above on all instances built on the Nitro system.

Type: String

Valid Values: `ipv4 | ipv6`

Required: No

KmsKeyId

The ID of the KMS key used to encrypt the disk in the cluster.

Type: String

Required: No

LogDeliveryConfigurations.LogDeliveryConfigurationRequest.N

Specifies the destination, format and type of the logs.

Type: Array of LogDeliveryConfigurationRequest objects

Required: No

MultiAZEnabled

A flag indicating if you have Multi-AZ enabled to enhance fault tolerance. For more information, see Minimizing Downtime: Multi-AZ.

Type: Boolean

Required: No

NetworkType

Must be either `ipv4` | `ipv6` | `dual_stack`. IPv6 is supported for workloads using Valkey 7.2 and above, Redis OSS engine version 6.2 to 7.1 and Memcached engine version 1.6.6 and above on all instances built on the Nitro system.

Type: String

Valid Values: `ipv4 | ipv6 | dual_stack`

Required: No

NodeGroupConfiguration.NodeGroupConfiguration.N

A list of node group (shard) configuration options. Each node group (shard) configuration has the following members: `PrimaryAvailabilityZone`, `ReplicaAvailabilityZones`, `ReplicaCount`, and `Slots`.

If you're creating a Valkey or Redis OSS (cluster mode disabled) or a Valkey or Redis OSS (cluster mode enabled) replication group, you can use this parameter to individually configure each node group (shard), or you can omit this parameter. However, it is required when seeding a Valkey or Redis OSS (cluster mode enabled) cluster from a S3 rdb file. You must configure each node group (shard) using this parameter because you must specify the slots for each node group.

Type: Array of NodeGroupConfiguration objects

Required: No

NotificationTopicArn

The Amazon Resource Name (ARN) of the Amazon Simple Notification Service (SNS) topic to which notifications are sent.

> **Note**
>
> The Amazon SNS topic owner must be the same as the cluster owner.

Type: String

Required: No

NumCacheClusters

The number of clusters this replication group initially has.

This parameter is not used if there is more than one node group (shard). You should use `ReplicasPerNodeGroup` instead.

If `AutomaticFailoverEnabled` is `true`, the value of this parameter must be at least 2. If `AutomaticFailoverEnabled` is `false` you can omit this parameter (it will default to 1), or you can explicitly set it to a value between 2 and 6.

The maximum permitted value for `NumCacheClusters` is 6 (1 primary plus 5 replicas).

Type: Integer

Required: No

NumNodeGroups

An optional parameter that specifies the number of node groups (shards) for this Valkey or Redis OSS (cluster mode enabled) replication group. For Valkey or Redis OSS (cluster mode disabled) either omit this parameter or set it to 1.

Default: 1

Type: Integer

Required: No

Port

The port number on which each member of the replication group accepts connections.

Type: Integer

Required: No

PreferredCacheClusterAZs.AvailabilityZone.N

A list of EC2 Availability Zones in which the replication group's clusters are created. The order of the Availability Zones in the list is the order in which clusters are allocated. The primary cluster is created in the first AZ in the list.

This parameter is not used if there is more than one node group (shard). You should use `NodeGroupConfiguration` instead.

> **Note**
>
> If you are creating your replication group in an Amazon VPC (recommended), you can only locate clusters in Availability Zones associated with the subnets in the selected subnet group.
> The number of Availability Zones listed must equal the value of `NumCacheClusters`.

Default: system chosen Availability Zones.

Type: Array of strings

Required: No

PreferredMaintenanceWindow

Specifies the weekly time range during which maintenance on the cluster is performed. It is specified as a range in the format ddd:hh24:mi-ddd:hh24:mi (24H Clock UTC). The minimum maintenance window is a 60 minute period.

Valid values for ddd are:

- `sun`
- `mon`
- `tue`
- `wed`
- `thu`
- `fri`
- `sat`

Example: `sun:23:00-mon:01:30`

Type: String

Required: No

PrimaryClusterId

The identifier of the cluster that serves as the primary for this replication group. This cluster must already exist and have a status of `available`.

This parameter is not required if `NumCacheClusters`, `NumNodeGroups`, or `ReplicasPerNodeGroup` is specified.

Type: String

Required: No

ReplicasPerNodeGroup

An optional parameter that specifies the number of replica nodes in each node group (shard). Valid values are 0 to 5.

Type: Integer

Required: No

SecurityGroupIds.SecurityGroupId.N

One or more Amazon VPC security groups associated with this replication group.

Use this parameter only when you are creating a replication group in an Amazon Virtual Private Cloud (Amazon VPC).

Type: Array of strings

Required: No

ServerlessCacheSnapshotName

The name of the snapshot used to create a replication group. Available for Valkey, Redis OSS only.

Type: String

Required: No

SnapshotArns.SnapshotArn.N

A list of Amazon Resource Names (ARN) that uniquely identify the Valkey or Redis OSS RDB snapshot files stored in Amazon S3. The snapshot files are used to populate the new replication group. The Amazon S3 object name in the ARN cannot contain any commas. The new replication group will have the number of node groups (console: shards) specified by the parameter *NumNodeGroups* or the number of node groups configured by *NodeGroupConfiguration* regardless of the number of ARNs specified here.

Example of an Amazon S3 ARN: `arn:aws:s3:::my_bucket/snapshot1.rdb`

Type: Array of strings

Required: No

SnapshotName

The name of a snapshot from which to restore data into the new replication group. The snapshot status changes to `restoring` while the new replication group is being created.

Type: String

Required: No

SnapshotRetentionLimit

The number of days for which ElastiCache retains automatic snapshots before deleting them. For example, if you set `SnapshotRetentionLimit` to 5, a snapshot that was taken today is retained for 5 days before being deleted.

Default: 0 (i.e., automatic backups are disabled for this cluster).

Type: Integer

Required: No

SnapshotWindow

The daily time range (in UTC) during which ElastiCache begins taking a daily snapshot of your node group (shard).

Example: `05:00-09:00`

If you do not specify this parameter, ElastiCache automatically chooses an appropriate time range.

Type: String

Required: No

Tags.Tag.N

A list of tags to be added to this resource. Tags are comma-separated key,value pairs (e.g. Key=`myKey`, Value=`myKeyValue`. You can include multiple tags as shown following: Key=`myKey`, Value=`myKeyValue` Key=`mySecondKey`, Value=`mySecondKeyValue`. Tags on replication groups will be replicated to all nodes.

Type: Array of Tag objects

Required: No

TransitEncryptionEnabled

A flag that enables in-transit encryption when set to `true`.

This parameter is valid only if the `Engine` parameter is `redis`, the `EngineVersion` parameter is `3.2.6`, `4.x` or later, and the cluster is being created in an Amazon VPC.

If you enable in-transit encryption, you must also specify a value for `CacheSubnetGroup`.

Required: Only available when creating a replication group in an Amazon VPC using Redis OSS version `3.2.6`, `4.x` or later.

Default: `false`

> **Important**
>
> For HIPAA compliance, you must specify `TransitEncryptionEnabled` as `true`, an `AuthToken`, and a `CacheSubnetGroup`.

Type: Boolean

Required: No

TransitEncryptionMode

A setting that allows you to migrate your clients to use in-transit encryption, with no downtime.

When setting `TransitEncryptionEnabled` to `true`, you can set your `TransitEncryptionMode` to `preferred` in the same request, to allow both encrypted and unencrypted connections at the same time. Once you migrate all your Valkey or Redis OSS clients to use encrypted connections you can modify the value to `required` to allow encrypted connections only.

Setting `TransitEncryptionMode` to `required` is a two-step process that requires you to first set the `TransitEncryptionMode` to `preferred`, after that you can set `TransitEncryptionMode` to `required`.

This process will not trigger the replacement of the replication group.

Type: String

Valid Values: `preferred | required`

Required: No

UserGroupIds.member.N

The user group to associate with the replication group.

Type: Array of strings

Array Members: Minimum number of 1 item.

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

Response Elements

The following element is returned by the service.

ReplicationGroup

Contains all of the attributes of a specific Valkey or Redis OSS replication group.

Type: ReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterNotFound

The requested cluster ID does not refer to an existing cluster.

HTTP Status Code: 404

CacheParameterGroupNotFound

The requested cache parameter group name does not refer to an existing cache parameter group.

HTTP Status Code: 404

CacheSecurityGroupNotFound

The requested cache security group name does not refer to an existing cache security group.

HTTP Status Code: 404

CacheSubnetGroupNotFoundFault

The requested cache subnet group name does not refer to an existing cache subnet group.

HTTP Status Code: 400

ClusterQuotaForCustomerExceeded

The request cannot be processed because it would exceed the allowed number of clusters per customer.

HTTP Status Code: 400

GlobalReplicationGroupNotFoundFault

The Global datastore does not exist

HTTP Status Code: 404

InsufficientCacheClusterCapacity

The requested cache node type is not available in the specified Availability Zone. For more information, see InsufficientCacheClusterCapacity in the ElastiCache User Guide.

HTTP Status Code: 400

InvalidCacheClusterState

The requested cluster is not in the `available` state.

HTTP Status Code: 400

InvalidGlobalReplicationGroupState

The Global datastore is not available or in primary-only state.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidUserGroupState

The user group is not in an active state.

HTTP Status Code: 400

InvalidVPCNetworkStateFault

The VPC network is in an invalid state.

HTTP Status Code: 400

NodeGroupsPerReplicationGroupQuotaExceeded

The request cannot be processed because it would exceed the maximum allowed number of node groups (shards) in a single replication group. The default maximum is 90

HTTP Status Code: 400

NodeQuotaForClusterExceeded

The request cannot be processed because it would exceed the allowed number of cache nodes in a single cluster.

HTTP Status Code: 400

NodeQuotaForCustomerExceeded

The request cannot be processed because it would exceed the allowed number of cache nodes per customer.

HTTP Status Code: 400

ReplicationGroupAlreadyExists

The specified replication group already exists.

HTTP Status Code: 400

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

UserGroupNotFound

The user group was not found or does not exist

HTTP Status Code: 404

Examples

CreateReplicationGroup - Valkey or Redis OSS (cluster mode disabled) Replication Group

The following example creates a Valkey or Redis OSS (cluster mode disabled) replication group with three nodes (`NumCacheClusters=3`), a primary and two read replicas. Because a single node group (shard) replication group technically could be either clustered or non-clustered, the parameter group `default.redis3.2` is specified, making this is a non-clustered replication group.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=CreateReplicationGroup
   &CacheParameterGroup=default.redis3.2
   &Engine=redis
   &EngineVersion=3.2.4
   &NumCacheClusters=3
   &ReplicationGroupDescription=My%20replication%20group
   &ReplicationGroupId=my-repgroup
   &PrimaryClusterId=my-redis-primary
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<CreateReplicationGroupResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
   <CreateReplicationGroupResult>
      <ReplicationGroup>
         <SnapshottingClusterId>my-redis-primary</SnapshottingClusterId>
```

```
         <MemberClusters>
            <ClusterId>my-redis-primary</ClusterId>
         </MemberClusters>
         <ReplicationGroupId>my-repgroup</ReplicationGroupId>
         <Status>creating</Status>
         <PendingModifiedValues />
         <Description>My replication group</Description>
      </ReplicationGroup>
   </CreateReplicationGroupResult>
   <ResponseMetadata>
      <RequestId>f3b7b32d-b9d2-11e3-8a16-7978bb24ffdf</RequestId>
   </ResponseMetadata>
</CreateReplicationGroupResponse>
```

Redis OSS (cluster mode enabled) Replication Group - all shards same profile

The following example creates a Valkey or Redis OSS (cluster mode enabled) replication group with three node groups (shards) and four replica nodes in each node group (shard). Note the following parameters and their values.

- EngineVersion=3.2.4
- CacheParameterGroup=default.redis3.2.cluster.on
- NumNodeGroups=3
- ReplicasPerNodeGroup=4

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=CreateReplicationGroup
   &CacheParameterGroup=default.redis3.2.cluster.on
   &Engine=redis
   &EngineVersion=3.2.4
   &ReplicationGroupDescription=My%20replication%20group
   &ReplicationGroupId=my-repgroup
   &NumNodeGroups=3
   &PrimaryClusterId=my-redis-primary
   &ReplicasPerNodeGroup=4
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
```

```
   &X-Amz-Credential=<credential>
```

Valkey or Redis OSS (cluster mode enabled) Replication Group - each shard configured separately

The following example creates a Valkey or Redis OSS (cluster mode enabled) replication group with two node groups (shards). The first shard has two replica nodes and slots 0-8192. The second shard has one replica and slots 8193-16383. Note the following parameters and their values.

- EngineVersion
- CacheParameterGroup
- NodeGroupConfiguration.NodeGroupConfiguration.*n*.PrimaryAvailabilityZone
- NodeGroupConfiguration.NodeGroupConfiguration.*n*.ReplicaAvailabilityZones.AvailabilityZone.*n*
- NodeGroupConfiguration.NodeGroupConfiguration.*n*.ReplicaCount
- NodeGroupConfiguration.NodeGroupConfiguration.*n*.Slots

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=CreateReplicationGroup
   &CacheParameterGroup=default.redis3.2.cluster.on
   &Engine=redis
   &EngineVersion=3.2.4
   &ReplicationGroupDescription=My%20replication%20group
   &ReplicationGroupId=my-repgroup
   &NodeGroupConfiguration.NodeGroupConfiguration.1.PrimaryAvailabilityZone=us-east-2a

 &NodeGroupConfiguration.NodeGroupConfiguration.1.ReplicaAvailabilityZones.AvailabilityZone.1=u
east-2b

 &NodeGroupConfiguration.NodeGroupConfiguration.1.ReplicaAvailabilityZones.AvailabilityZone.2=u
east-2c
   &NodeGroupConfiguration.NodeGroupConfiguration.1.ReplicaCount=2
   &NodeGroupConfiguration.NodeGroupConfiguration.1.Slots=0-8192
   &NodeGroupConfiguration.NodeGroupConfiguration.2.PrimaryAvailabilityZone=us-east-2b

 &NodeGroupConfiguration.NodeGroupConfiguration.2.ReplicaAvailabilityZones.AvailabilityZone.1=u
east-2d
   &NodeGroupConfiguration.NodeGroupConfiguration.2.ReplicaCount=1
```

```
    &NodeGroupConfiguration.NodeGroupConfiguration.2.Slots=8193-16383
    &PrimaryClusterId=my-redis-primary
    &Version=2015-02-02
    &SignatureVersion=4
    &SignatureMethod=HmacSHA256
    &Timestamp=20150202T192317Z
    &X-Amz-Credential=<credential>
```

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

CreateServerlessCache

Creates a serverless cache.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

Engine

The name of the cache engine to be used for creating the serverless cache.

Type: String

Required: Yes

ServerlessCacheName

User-provided identifier for the serverless cache. This parameter is stored as a lowercase string.

Type: String

Required: Yes

CacheUsageLimits

Sets the cache usage limits for storage and ElastiCache Processing Units for the cache.

Type: CacheUsageLimits object

Required: No

DailySnapshotTime

The daily time that snapshots will be created from the new serverless cache. By default this number is populated with 0, i.e. no snapshots will be created on an automatic daily basis. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: No

Description

User-provided description for the serverless cache. The default is NULL, i.e. if no description is provided then an empty string will be returned. The maximum length is 255 characters.

Type: String

Required: No

KmsKeyId

ARN of the customer managed key for encrypting the data at rest. If no KMS key is provided, a default service key is used.

Type: String

Required: No

MajorEngineVersion

The version of the cache engine that will be used to create the serverless cache.

Type: String

Required: No

SecurityGroupIds.SecurityGroupId.N

A list of the one or more VPC security groups to be associated with the serverless cache. The security group will authorize traffic access for the VPC end-point (private-link). If no other information is given this will be the VPC's Default Security Group that is associated with the cluster VPC end-point.

Type: Array of strings

Required: No

SnapshotArnsToRestore.SnapshotArn.N

The ARN(s) of the snapshot that the new serverless cache will be created from. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: Array of strings

Required: No

SnapshotRetentionLimit

The number of snapshots that will be retained for the serverless cache that is being created. As new snapshots beyond this limit are added, the oldest snapshots will be deleted on a rolling basis. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: Integer

Required: No

SubnetIds.SubnetId.N

A list of the identifiers of the subnets where the VPC endpoint for the serverless cache will be deployed. All the subnetIds must belong to the same VPC.

Type: Array of strings

Required: No

Tags.Tag.N

The list of tags (key, value) pairs to be added to the serverless cache resource. Default is NULL.

Type: Array of Tag objects

Required: No

UserGroupId

The identifier of the UserGroup to be associated with the serverless cache. Available for Valkey and Redis OSS only. Default is NULL.

Type: String

Required: No

Response Elements

The following element is returned by the service.

ServerlessCache

The response for the attempt to create the serverless cache.

Type: ServerlessCache object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidCredentials

You must enter valid credentials.

HTTP Status Code: 408

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidServerlessCacheStateFault

The account for these credentials is not currently active.

HTTP Status Code: 400

InvalidUserGroupState

The user group is not in an active state.

HTTP Status Code: 400

ServerlessCacheAlreadyExistsFault

A serverless cache with this name already exists.

HTTP Status Code: 400

ServerlessCacheNotFoundFault

The serverless cache was not found or does not exist.

HTTP Status Code: 404

ServerlessCacheQuotaForCustomerExceededFault

The number of serverless caches exceeds the customer quota.

HTTP Status Code: 400

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

UserGroupNotFound

The user group was not found or does not exist

HTTP Status Code: 404

Examples

CreateServerlessCache with Redis OSS

This example illustrates one usage of CreateServerlessCache.

Sample Request

```
{
    "input": {
        "ServerlessCacheName": "my-serverless-cache",
        "Description": "A serverless cache.",
        "Engine": "redis",
        "MajorEngineVersion": "7",
        "SubnetIds": [
            "subnet-xxx8c982",
            "subnet-xxx382f3",
            "subnet-xxxb3e7c0"
        ],
        "CacheUsageLimits": {
            "DataStorage" : {
                "Maximum" : 10,
                "Unit" : "GB"
            },
            "ECPUPerSecond" : {
                "Maximum" : 50000
            }
```

```
        },
        "SecurityGroupIds": [
            "sg-xxx0c9af"
        ],
        "SnapshotRetentionLimit": 10,
        "DailySnapshotTime": "09:00"
    },
    "output": {
        "ServerlessCache": {
            "ServerlessCacheName": "my-serverless-cache",
            "Description": "A serverless cache.",
            "Status": "creating",
            "Engine": "redis",
            "MajorEngineVersion": "7",
            "FullEngineVersion": "7.0",
            "SubnetIds": [
                "subnet-xxx8c982",
                "subnet-xxx382f3",
                "subnet-xxxb3e7c0"
            ],
            "CacheUsageLimits": {
                "DataStorage" : {
                    "Maximum" : 10,
                    "Unit" : "GB"
                },
                "ECPUPerSecond" : {
                    "Maximum" : 50000
                }
            },
            "SecurityGroupIds": [
                "sg-xxx0c9af"
            ],
            "ARN": "arn:aws:elasticache:us-east-1:222222222222:serverlesscache:my-
serverless-cache",
            "SnapshotRetentionLimit": 10,
            "DailySnapshotTime": "09:00"
        }
    }
}
```

CreateServerlessCache (Memcached)

This example illustrates one usage of CreateServerlessCache.

Sample Request

```
{
    "input": {
        "ServerlessCacheName": "my-serverless-cache",
        "Description": "A serverless cache.",
        "Engine": "memcached",
        "MajorEngineVersion": "1.6",
        "SubnetIds": [
            "subnet-xxx8c982",
            "subnet-xxx382f3",
            "subnet-xxxb3e7c0"
        ],
        "CacheUsageLimits": {
            "DataStorage" : {
                "Maximum" : 10,
                "Unit" : "GB"
            },
            "ECPUPerSecond" : {
                "Maximum" : 50000
            }
        },
        "SecurityGroupIds": [
            "sg-xxx0c9af"
        ],
            "ARN": "arn:aws:elasticache:us-east-1:222222222222:serverlesscache:my-
serverless-cache",
            "SnapshotRetentionLimit": 10,
            "DailySnapshotTime": "09:00"
    },
    "output": {
        "ServerlessCache": {
            "ServerlessCacheName": "my-serverless-cache",
            "Description": "A serverless cache.",
            "Status": "creating",
            "Engine": "memcached",
            "MajorEngineVersion": "1.6",
            "FullEngineVersion": "1.6.21",
            "SubnetIds": [
                "subnet-xxx8c982",
                "subnet-xxx382f3",
                "subnet-xxxb3e7c0"
            ],
            "CacheUsageLimits": {
```

```
                "DataStorage" : {
                    "Maximum" : 10,
                    "Unit" : "GB"
                },
                "ECPUPerSecond" : {
                    "Maximum" : 50000
                }
            },
            "SecurityGroupIds": [
                "sg-xxx0c9af"
            ],
            "ARN": "arn:aws:elasticache:us-east-1:222222222222:serverlesscache:my-
serverless-cache",
        }
    }
}
```

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

CreateServerlessCacheSnapshot

This API creates a copy of an entire ServerlessCache at a specific moment in time. Available for Valkey, Redis OSS and Serverless Memcached only.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ServerlessCacheName

The name of an existing serverless cache. The snapshot is created from this cache. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: Yes

ServerlessCacheSnapshotName

The name for the snapshot being created. Must be unique for the customer account. Available for Valkey, Redis OSS and Serverless Memcached only. Must be between 1 and 255 characters.

Type: String

Required: Yes

KmsKeyId

The ID of the KMS key used to encrypt the snapshot. Available for Valkey, Redis OSS and Serverless Memcached only. Default: NULL

Type: String

Required: No

Tags.Tag.N

A list of tags to be added to the snapshot resource. A tag is a key-value pair. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: Array of Tag objects

Required: No

Response Elements

The following element is returned by the service.

ServerlessCacheSnapshot

The state of a serverless cache snapshot at a specific point in time, to the millisecond. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: ServerlessCacheSnapshot object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidServerlessCacheStateFault

The account for these credentials is not currently active.

HTTP Status Code: 400

ServerlessCacheNotFoundFault

The serverless cache was not found or does not exist.

HTTP Status Code: 404

ServerlessCacheSnapshotAlreadyExistsFault

A serverless cache snapshot with this name already exists. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 400

ServerlessCacheSnapshotQuotaExceededFault

The number of serverless cache snapshots exceeds the customer snapshot quota. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 400

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

Examples

CreateServerlessCacheSnapshot

This example illustrates one usage of CreateServerlessCacheSnapshot.

Sample Request

```
{
    "input": {
        "ServerlessCacheSnapshotName": "my-serverless-cache-snapshot",
        "ServerlessCacheName": "my-serverless-cache"
    },
    "output": {
        "ServerlessCacheSnapshot": {
            "ServerlessCacheSnapshotName": "my-serverless-cache-snapshot",
            "ARN": "arn:aws:elasticache:us-
east-1:222222222222:serverlesscachesnapshot:my-serverless-cache-snapshot",
            "SnapshotType": "manual",
            "Status": "creating",
            "ServerlessCacheConfiguration": {
                "ServerlessCacheName": "my-serverless-cache",
                "Engine": "redis",
```

```
                "MajorEngineVersion": "7"
            }
        }
    }
}
```

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

CreateSnapshot

Creates a copy of an entire cluster or replication group at a specific moment in time.

> **Note**
>
> This operation is valid for Valkey or Redis OSS only.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

SnapshotName

A name for the snapshot being created.

Type: String

Required: Yes

CacheClusterId

The identifier of an existing cluster. The snapshot is created from this cluster.

Type: String

Required: No

KmsKeyId

The ID of the KMS key used to encrypt the snapshot.

Type: String

Required: No

ReplicationGroupId

The identifier of an existing replication group. The snapshot is created from this replication group.

Type: String

Required: No

Tags.Tag.N

A list of tags to be added to this resource. A tag is a key-value pair. A tag key must be accompanied by a tag value, although null is accepted.

Type: Array of Tag objects

Required: No

Response Elements

The following element is returned by the service.

Snapshot

Represents a copy of an entire Valkey or Redis OSS cluster as of the time when the snapshot was taken.

Type: Snapshot object

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterNotFound

The requested cluster ID does not refer to an existing cluster.

HTTP Status Code: 404

InvalidCacheClusterState

The requested cluster is not in the `available` state.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

SnapshotAlreadyExistsFault

You already have a snapshot with the given name.

HTTP Status Code: 400

SnapshotFeatureNotSupportedFault

You attempted one of the following operations:

- Creating a snapshot of a Valkey or Redis OSS cluster running on a `cache.t1.micro` cache node.
- Creating a snapshot of a cluster that is running Memcached rather than Valkey or Redis OSS.

Neither of these are supported by ElastiCache.

HTTP Status Code: 400

SnapshotQuotaExceededFault

The request cannot be processed because it would exceed the maximum number of snapshots.

HTTP Status Code: 400

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

Examples

CreateSnapshot

This example illustrates one usage of CreateSnapshot.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=CreateSnapshot
   &CacheClusterId=my-redis-primary
   &SnapshotName=my-manual-snapshot
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<CreateSnapshotResponse xmlns="http://elasticache.amazonaws.com/doc/2015-02-02/">
   <CreateSnapshotResult>
      <Snapshot>
         <CacheClusterId>my-redis-primary</CacheClusterId>
         <Port>6379</Port>
         <CacheNodeType>cache.m1.small</CacheNodeType>
         <CacheParameterGroupName>default.redis2.8</CacheParameterGroupName>
         <Engine>redis</Engine>
         <PreferredAvailabilityZone>us-west-2c</PreferredAvailabilityZone>
         <CacheClusterCreateTime>2015-02-02T18:46:57.972Z</CacheClusterCreateTime>
         <EngineVersion>2.8.6</EngineVersion>
         <SnapshotSource>manual</SnapshotSource>
         <AutoMinorVersionUpgrade>true</AutoMinorVersionUpgrade>
         <PreferredMaintenanceWindow>wed:09:00-wed:10:00</PreferredMaintenanceWindow>
         <SnapshotName>my-manual-snapshot</SnapshotName>
         <SnapshotRetentionLimit>5</SnapshotRetentionLimit>
         <NodeSnapshots>
            <NodeSnapshot>
               <CacheNodeCreateTime>2015-02-02T18:46:57.972Z</CacheNodeCreateTime>
               <CacheNodeId>0001</CacheNodeId>
```

```
                <CacheSize />
             </NodeSnapshot>
          </NodeSnapshots>
          <SnapshotStatus>creating</SnapshotStatus>
          <NumCacheNodes>1</NumCacheNodes>
          <SnapshotWindow>07:30-08:30</SnapshotWindow>
       </Snapshot>
    </CreateSnapshotResult>
    <ResponseMetadata>
       <RequestId>faf5a232-b9ce-11e3-8a16-7978bb24ffdf</RequestId>
    </ResponseMetadata>
</CreateSnapshotResponse>
```

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

CreateUser

For Valkey engine version 7.2 onwards and Redis OSS 6.0 to 7.1: Creates a user. For more information, see Using Role Based Access Control (RBAC).

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

AccessString

Access permissions string used for this user.

Type: String

Pattern: `.*\S.*`

Required: Yes

Engine

The options are valkey or redis.

Type: String

Pattern: `[a-zA-Z]*`

Required: Yes

UserId

The ID of the user.

Type: String

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: Yes

UserName

The username of the user.

Type: String

Length Constraints: Minimum length of 1.

Required: Yes

AuthenticationMode

Specifies how to authenticate the user.

Type: AuthenticationMode object

Required: No

NoPasswordRequired

Indicates a password is not required for this user.

Type: Boolean

Required: No

Passwords.member.N

Passwords used for this user. You can create up to two passwords for each user.

Type: Array of strings

Array Members: Minimum number of 1 item.

Required: No

Tags.Tag.N

A list of tags to be added to this resource. A tag is a key-value pair. A tag key must be accompanied by a tag value, although null is accepted.

Type: Array of Tag objects

Required: No

Response Elements

The following elements are returned by the service.

AccessString

Access permissions string used for this user.

Type: String

ARN

The Amazon Resource Name (ARN) of the user.

Type: String

Authentication

Denotes whether the user requires a password to authenticate.

Type: Authentication object

Engine

The options are valkey or redis.

Type: String

Pattern: `[a-zA-Z]*`

MinimumEngineVersion

The minimum engine version required, which is Redis OSS 6.0

Type: String

Status

Indicates the user status. Can be "active", "modifying" or "deleting".

Type: String

UserGroupIds.member.N

Returns a list of the user group IDs the user belongs to.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

UserId

The ID of the user.

Type: String

UserName

The username of the user.

Type: String

Errors

For information about the errors that are common to all actions, see Common Errors.

DuplicateUserName

A user with this username already exists.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

UserAlreadyExists

A user with this ID already exists.

HTTP Status Code: 400

UserQuotaExceeded

The quota of users has been exceeded.

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

CreateUserGroup

For Valkey engine version 7.2 onwards and Redis OSS 6.0 to 7.1: Creates a user group. For more information, see Using Role Based Access Control (RBAC)

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

Engine

Sets the engine listed in a user group. The options are valkey or redis.

Type: String

Pattern: `[a-zA-Z]*`

Required: Yes

UserGroupId

The ID of the user group.

Type: String

Required: Yes

Tags.Tag.N

A list of tags to be added to this resource. A tag is a key-value pair. A tag key must be accompanied by a tag value, although null is accepted. Available for Valkey and Redis OSS only.

Type: Array of Tag objects

Required: No

UserIds.member.N

The list of user IDs that belong to the user group.

Type: Array of strings

Array Members: Minimum number of 1 item.

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

Response Elements

The following elements are returned by the service.

ARN

The Amazon Resource Name (ARN) of the user group.

Type: String

Engine

The options are valkey or redis.

Type: String

Pattern: `[a-zA-Z]*`

MinimumEngineVersion

The minimum engine version required, which is Redis OSS 6.0

Type: String

PendingChanges

A list of updates being applied to the user group.

Type: UserGroupPendingChanges object

ReplicationGroups.member.N

A list of replication groups that the user group can access.

Type: Array of strings

ServerlessCaches.member.N

Indicates which serverless caches the specified user group is associated with. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: Array of strings

Status

Indicates user group status. Can be "creating", "active", "modifying", "deleting".

Type: String

UserGroupId

The ID of the user group.

Type: String

UserIds.member.N

The list of user IDs that belong to the user group.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Errors

For information about the errors that are common to all actions, see Common Errors.

DefaultUserRequired

You must add default user to a user group.

HTTP Status Code: 400

DuplicateUserName

A user with this username already exists.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

UserGroupAlreadyExists

The user group with this ID already exists.

HTTP Status Code: 400

UserGroupQuotaExceeded

The number of users exceeds the user group limit.

HTTP Status Code: 400

UserNotFound

The user does not exist or could not be found.

HTTP Status Code: 404

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

DecreaseNodeGroupsInGlobalReplicationGroup

Decreases the number of node groups in a Global datastore

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ApplyImmediately

Indicates that the shard reconfiguration process begins immediately. At present, the only permitted value for this parameter is true.

Type: Boolean

Required: Yes

GlobalReplicationGroupId

The name of the Global datastore

Type: String

Required: Yes

NodeGroupCount

The number of node groups (shards) that results from the modification of the shard configuration

Type: Integer

Required: Yes

GlobalNodeGroupsToRemove.GlobalNodeGroupId.N

If the value of NodeGroupCount is less than the current number of node groups (shards), then either NodeGroupsToRemove or NodeGroupsToRetain is required. GlobalNodeGroupsToRemove is a list of NodeGroupIds to remove from the cluster. ElastiCache will attempt to remove all node groups listed by GlobalNodeGroupsToRemove from the cluster.

Type: Array of strings

Required: No

GlobalNodeGroupsToRetain.GlobalNodeGroupId.N

If the value of NodeGroupCount is less than the current number of node groups (shards), then either NodeGroupsToRemove or NodeGroupsToRetain is required. GlobalNodeGroupsToRetain is a list of NodeGroupIds to retain from the cluster. ElastiCache will attempt to retain all node groups listed by GlobalNodeGroupsToRetain from the cluster.

Type: Array of strings

Required: No

Response Elements

The following element is returned by the service.

GlobalReplicationGroup

Consists of a primary cluster that accepts writes and an associated secondary cluster that resides in a different Amazon region. The secondary cluster accepts only reads. The primary cluster automatically replicates updates to the secondary cluster.

- The **GlobalReplicationGroupIdSuffix** represents the name of the Global datastore, which is what you use to associate a secondary cluster.

Type: GlobalReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

GlobalReplicationGroupNotFoundFault

The Global datastore does not exist

HTTP Status Code: 404

InvalidGlobalReplicationGroupState

The Global datastore is not available or in primary-only state.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

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

DecreaseReplicaCount

Dynamically decreases the number of replicas in a Valkey or Redis OSS (cluster mode disabled) replication group or the number of replica nodes in one or more node groups (shards) of a Valkey or Redis OSS (cluster mode enabled) replication group. This operation is performed with no cluster down time.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ApplyImmediately

If `True`, the number of replica nodes is decreased immediately. `ApplyImmediately=False` is not currently supported.

Type: Boolean

Required: Yes

ReplicationGroupId

The id of the replication group from which you want to remove replica nodes.

Type: String

Required: Yes

NewReplicaCount

The number of read replica nodes you want at the completion of this operation. For Valkey or Redis OSS (cluster mode disabled) replication groups, this is the number of replica nodes in the replication group. For Valkey or Redis OSS (cluster mode enabled) replication groups, this is the number of replica nodes in each of the replication group's node groups.

The minimum number of replicas in a shard or replication group is:

- Valkey or Redis OSS (cluster mode disabled)
  - If Multi-AZ is enabled: 1
  - If Multi-AZ is not enabled: 0
- Valkey or Redis OSS (cluster mode enabled): 0 (though you will not be able to failover to a replica if your primary node fails)

Type: Integer

Required: No

ReplicaConfiguration.ConfigureShard.N

A list of `ConfigureShard` objects that can be used to configure each shard in a Valkey or Redis OSS replication group. The `ConfigureShard` has three members: `NewReplicaCount`, `NodeGroupId`, and `PreferredAvailabilityZones`.

Type: Array of ConfigureShard objects

Required: No

ReplicasToRemove.member.N

A list of the node ids to remove from the replication group or node group (shard).

Type: Array of strings

Required: No

Response Elements

The following element is returned by the service.

ReplicationGroup

Contains all of the attributes of a specific Valkey or Redis OSS replication group.

Type: ReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

ClusterQuotaForCustomerExceeded

The request cannot be processed because it would exceed the allowed number of clusters per customer.

HTTP Status Code: 400

InsufficientCacheClusterCapacity

The requested cache node type is not available in the specified Availability Zone. For more information, see InsufficientCacheClusterCapacity in the ElastiCache User Guide.

HTTP Status Code: 400

InvalidCacheClusterState

The requested cluster is not in the `available` state.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

InvalidVPCNetworkStateFault

The VPC network is in an invalid state.

HTTP Status Code: 400

NodeGroupsPerReplicationGroupQuotaExceeded

The request cannot be processed because it would exceed the maximum allowed number of node groups (shards) in a single replication group. The default maximum is 90

HTTP Status Code: 400

NodeQuotaForCustomerExceeded

The request cannot be processed because it would exceed the allowed number of cache nodes per customer.

HTTP Status Code: 400

NoOperationFault

The operation was not performed because no changes were required.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

Examples

Example

The following example removes two replicas from each node group in the replication group `sample-repl-group`.

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DecreaseReplicaCount
   &ApplyImmediately=True
   &NewReplicaCount=2
   &ReplicasToRemove.ReplicaToRemove.1=0001
   &ReplicasToRemove.ReplicaToRemove.2=0003
   &ReplicationGroupId=sample-repl-group
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Example

The following example removes replicas from two node groups. Because there are multiple node groups, this example is for a Valkey or Redis OSS (cluster mode enabled) replication group.

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DecreaseReplicaCount
   &ApplyImmediately=True
   &ReplicaConfiguration.ConfigureShard.1.NodeGroupId=0001
   &ReplicaConfiguration.ConfigureShard.1.NewReplicaCount=1

 &ReplicaConfiguration.ConfigureShard.1.PreferredAvailabilityZones.PreferredAvailabilityZone.1=
east-1a

 &ReplicaConfiguration.ConfigureShard.1.PreferredAvailabilityZones.PreferredAvailabilityZone.2=
east-1c
   &ReplicaConfiguration.ConfigureShard.2.NodeGroupId=0003
   &ReplicaConfiguration.ConfigureShard.2.NewReplicaCount=2

 &ReplicaConfiguration.ConfigureShard.2.PreferredAvailabilityZones.PreferredAvailabilityZone.1=
east-1a

 &ReplicaConfiguration.ConfigureShard.2.PreferredAvailabilityZones.PreferredAvailabilityZone.2=
east-1b

 &ReplicaConfiguration.ConfigureShard.2.PreferredAvailabilityZones.PreferredAvailabilityZone.4=
east-1c
   &ReplicationGroupId=samplem--repl-group
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

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

DeleteCacheCluster

Deletes a previously provisioned cluster. `DeleteCacheCluster` deletes all associated cache nodes, node endpoints and the cluster itself. When you receive a successful response from this operation, Amazon ElastiCache immediately begins deleting the cluster; you cannot cancel or revert this operation.

This operation is not valid for:

- Valkey or Redis OSS (cluster mode enabled) clusters
- Valkey or Redis OSS (cluster mode disabled) clusters
- A cluster that is the last read replica of a replication group
- A cluster that is the primary node of a replication group
- A node group (shard) that has Multi-AZ mode enabled
- A cluster from a Valkey or Redis OSS (cluster mode enabled) replication group
- A cluster that is not in the `available` state

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheClusterId

The cluster identifier for the cluster to be deleted. This parameter is not case sensitive.

Type: String

Required: Yes

FinalSnapshotIdentifier

The user-supplied name of a final cluster snapshot. This is the unique name that identifies the snapshot. ElastiCache creates the snapshot, and then deletes the cluster immediately afterward.

Type: String

Required: No

Response Elements

The following element is returned by the service.

CacheCluster

Contains all of the attributes of a specific cluster.

Type: CacheCluster object

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterNotFound

The requested cluster ID does not refer to an existing cluster.

HTTP Status Code: 404

InvalidCacheClusterState

The requested cluster is not in the `available` state.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

SnapshotAlreadyExistsFault

You already have a snapshot with the given name.

HTTP Status Code: 400

SnapshotFeatureNotSupportedFault

You attempted one of the following operations:

- Creating a snapshot of a Valkey or Redis OSS cluster running on a `cache.t1.micro` cache node.
- Creating a snapshot of a cluster that is running Memcached rather than Valkey or Redis OSS.

Neither of these are supported by ElastiCache.

HTTP Status Code: 400

SnapshotQuotaExceededFault

The request cannot be processed because it would exceed the maximum number of snapshots.

HTTP Status Code: 400

Examples

DeleteCacheCluster

This example illustrates one usage of DeleteCacheCluster.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DeleteCacheCluster
   &CacheClusterId=simcoprod43
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DeleteCacheClusterResponse xmlns="http://elasticache.amazonaws.com/doc/2015-02-02/">
   <DeleteCacheClusterResult>
      <CacheCluster>
         <CacheParameterGroup>
```

```
            <ParameterApplyStatus>in-sync</ParameterApplyStatus>
            <CacheParameterGroupName>default.memcached1.4</CacheParameterGroupName>
            <CacheNodeIdsToReboot/>
         </CacheParameterGroup>
         <CacheClusterId>simcoprod43</CacheClusterId>
         <CacheClusterStatus>deleting</CacheClusterStatus>
         <ConfigurationEndpoint>
            <Port>11211</Port>
            <Address>simcoprod43.m2st2p.cfg.cache.amazonaws.com</Address>
         </ConfigurationEndpoint>
         <CacheNodeType>cache.m1.large</CacheNodeType>
         <Engine>memcached</Engine>
         <PendingModifiedValues/>
         <PreferredAvailabilityZone>us-west-2b</PreferredAvailabilityZone>
         <CacheClusterCreateTime>2015-02-02T02:18:26.497Z</CacheClusterCreateTime>
         <EngineVersion>1.4.5</EngineVersion>
         <AutoMinorVersionUpgrade>true</AutoMinorVersionUpgrade>
         <PreferredMaintenanceWindow>mon:05:00-mon:06:00</PreferredMaintenanceWindow>
         <CacheSecurityGroups>
            <CacheSecurityGroup>
               <CacheSecurityGroupName>default</CacheSecurityGroupName>
               <Status>active</Status>
            </CacheSecurityGroup>
         </CacheSecurityGroups>
         <NumCacheNodes>3</NumCacheNodes>
      </CacheCluster>
   </DeleteCacheClusterResult>
   <ResponseMetadata>
      <RequestId>ab84aa7e-b7fa-11e0-9b0b-a9261be2b354</RequestId>
   </ResponseMetadata>
</DeleteCacheClusterResponse>
```


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

DeleteCacheParameterGroup

Deletes the specified cache parameter group. You cannot delete a cache parameter group if it is associated with any cache clusters. You cannot delete the default cache parameter groups in your account.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheParameterGroupName

The name of the cache parameter group to delete.

> **Note**
> The specified cache security group must not be associated with any clusters.

Type: String

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheParameterGroupNotFound

The requested cache parameter group name does not refer to an existing cache parameter group.

HTTP Status Code: 404

InvalidCacheParameterGroupState

The current state of the cache parameter group does not allow the requested operation to occur.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

Examples

DeleteCacheParameterGroup

This example illustrates one usage of DeleteCacheParameterGroup.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DeleteCacheParameterGroup
   &CacheParameterGroupName=myparametergroup
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DeleteCacheParameterGroupResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
  <ResponseMetadata>
    <RequestId>d0a417cb-575b-11e0-8869-cd22b4f9d96f</RequestId>
  </ResponseMetadata>
</DeleteCacheParameterGroupResponse>
```

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

DeleteCacheSecurityGroup

Deletes a cache security group.

> **Note**
> You cannot delete a cache security group if it is associated with any clusters.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheSecurityGroupName

The name of the cache security group to delete.

> **Note**
> You cannot delete the default security group.

Type: String

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheSecurityGroupNotFound

The requested cache security group name does not refer to an existing cache security group.

HTTP Status Code: 404

InvalidCacheSecurityGroupState

The current state of the cache security group does not allow deletion.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

Examples

DeleteCacheSecurityGroup

This example illustrates one usage of DeleteCacheSecurityGroup.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
?Action=DeleteCacheSecurityGroup
&CacheSecurityGroupName=mycachesecuritygroup3
&Version=2015-02-02
&SignatureVersion=4
&SignatureMethod=HmacSHA256
&Timestamp=20150202T192317Z
&X-Amz-Credential=<credential>
```

Sample Response

```
<DeleteCacheSecurityGroupResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
    <ResponseMetadata>
        <RequestId>c130cfb7-3650-11e0-ae57-f96cfe56749c</RequestId>
    </ResponseMetadata>
</DeleteCacheSecurityGroupResponse>
```

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

DeleteCacheSubnetGroup

Deletes a cache subnet group.

> **Note**
>
> You cannot delete a default cache subnet group or one that is associated with any clusters.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheSubnetGroupName

The name of the cache subnet group to delete.

Constraints: Must contain no more than 255 alphanumeric characters or hyphens.

Type: String

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheSubnetGroupInUse

The requested cache subnet group is currently in use.

HTTP Status Code: 400

CacheSubnetGroupNotFoundFault

The requested cache subnet group name does not refer to an existing cache subnet group.

HTTP Status Code: 400

Examples

DeleteCacheSubnetGroup

This example illustrates one usage of DeleteCacheSubnetGroup.

Sample Request

```
https://elasticache.amazonaws.com/
   ?Action=DeleteCacheSubnetGroup
   &CacheSubnetGroupName=mysubnetgroup
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DeleteCacheSubnetGroupResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
      <ResponseMetadata>
        <RequestId>5d013245-4172-11df-8520-e7e1e602a915</RequestId>
      </ResponseMetadata>
    </DeleteCacheSubnetGroupResponse>
```

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

DeleteGlobalReplicationGroup

Deleting a Global datastore is a two-step process:

- First, you must DisassociateGlobalReplicationGroup to remove the secondary clusters in the Global datastore.
- Once the Global datastore contains only the primary cluster, you can use the `DeleteGlobalReplicationGroup` API to delete the Global datastore while retainining the primary cluster using `RetainPrimaryReplicationGroup=true`.

Since the Global Datastore has only a primary cluster, you can delete the Global Datastore while retaining the primary by setting `RetainPrimaryReplicationGroup=true`. The primary cluster is never deleted when deleting a Global Datastore. It can only be deleted when it no longer is associated with any Global Datastore.

When you receive a successful response from this operation, Amazon ElastiCache immediately begins deleting the selected resources; you cannot cancel or revert this operation.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

GlobalReplicationGroupId

The name of the Global datastore

Type: String

Required: Yes

RetainPrimaryReplicationGroup

The primary replication group is retained as a standalone replication group.

Type: Boolean

Required: Yes

Response Elements

The following element is returned by the service.

GlobalReplicationGroup

Consists of a primary cluster that accepts writes and an associated secondary cluster that resides in a different Amazon region. The secondary cluster accepts only reads. The primary cluster automatically replicates updates to the secondary cluster.

- The **GlobalReplicationGroupIdSuffix** represents the name of the Global datastore, which is what you use to associate a secondary cluster.

Type: GlobalReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

GlobalReplicationGroupNotFoundFault

The Global datastore does not exist

HTTP Status Code: 404

InvalidGlobalReplicationGroupState

The Global datastore is not available or in primary-only state.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

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

DeleteReplicationGroup

Deletes an existing replication group. By default, this operation deletes the entire replication group, including the primary/primaries and all of the read replicas. If the replication group has only one primary, you can optionally delete only the read replicas, while retaining the primary by setting `RetainPrimaryCluster=true`.

When you receive a successful response from this operation, Amazon ElastiCache immediately begins deleting the selected resources; you cannot cancel or revert this operation.

> **Note**
>
> - `CreateSnapshot` permission is required to create a final snapshot. Without this permission, the API call will fail with an `Access Denied` exception.
> - This operation is valid for Redis OSS only.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ReplicationGroupId

The identifier for the cluster to be deleted. This parameter is not case sensitive.

Type: String

Required: Yes

FinalSnapshotIdentifier

The name of a final node group (shard) snapshot. ElastiCache creates the snapshot from the primary node in the cluster, rather than one of the replicas; this is to ensure that it captures the freshest data. After the final snapshot is taken, the replication group is immediately deleted.

Type: String

Required: No

RetainPrimaryCluster

If set to `true`, all of the read replicas are deleted, but the primary node is retained.

Type: Boolean

Required: No

Response Elements

The following element is returned by the service.

ReplicationGroup

Contains all of the attributes of a specific Valkey or Redis OSS replication group.

Type: ReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

SnapshotAlreadyExistsFault

You already have a snapshot with the given name.

HTTP Status Code: 400

SnapshotFeatureNotSupportedFault

You attempted one of the following operations:

- Creating a snapshot of a Valkey or Redis OSS cluster running on a `cache.t1.micro` cache node.
- Creating a snapshot of a cluster that is running Memcached rather than Valkey or Redis OSS.

Neither of these are supported by ElastiCache.

HTTP Status Code: 400

SnapshotQuotaExceededFault

The request cannot be processed because it would exceed the maximum number of snapshots.

HTTP Status Code: 400

Examples

DeleteReplicationGroup

This example illustrates one usage of DeleteReplicationGroup.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DeleteReplicationGroup &RetainPrimaryCluster=false
   &FinalSnapshotIdentifier=my-final-snapshot
   &ReplicationGroupId=my-repgroup
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DeleteReplicationGroupResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
   <DeleteReplicationGroupResult>
```

```
      <ReplicationGroup>
         <SnapshottingClusterId>my-redis-primary</SnapshottingClusterId>
decrease-replica-count         <ReplicationGroupId>my-repgroup</ReplicationGroupId>
         <Status>deleting</Status>
         <PendingModifiedValues />
         <Description>My replication group</Description>
      </ReplicationGroup>
   </DeleteReplicationGroupResult>
   <ResponseMetadata>
      <RequestId>93eb37db-b9d7-11e3-8a16-7978bb24ffdf</RequestId>
   </ResponseMetadata>
</DeleteReplicationGroupResponse>
```

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

DeleteServerlessCache

Deletes a specified existing serverless cache.

> **Note**
>
> `CreateServerlessCacheSnapshot` permission is required to create a final snapshot. Without this permission, the API call will fail with an `Access Denied` exception.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ServerlessCacheName

The identifier of the serverless cache to be deleted.

Type: String

Required: Yes

FinalSnapshotName

Name of the final snapshot to be taken before the serverless cache is deleted. Available for Valkey, Redis OSS and Serverless Memcached only. Default: NULL, i.e. a final snapshot is not taken.

Type: String

Required: No

Response Elements

The following element is returned by the service.

ServerlessCache

Provides the details of the specified serverless cache that is about to be deleted.

Type: ServerlessCache object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidCredentials

You must enter valid credentials.

HTTP Status Code: 408

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidServerlessCacheStateFault

The account for these credentials is not currently active.

HTTP Status Code: 400

ServerlessCacheNotFoundFault

The serverless cache was not found or does not exist.

HTTP Status Code: 404

ServerlessCacheSnapshotAlreadyExistsFault

A serverless cache snapshot with this name already exists. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 400

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

Examples

DeleteServerlessCache

This example illustrates one usage of DeleteServerlessCache.

Sample Request

```
{
    "input": {
        "ServerlessCacheName": "my-serverless-cache"
    },
    "output": {
        "ServerlessCache": {
            "ServerlessCacheName": "my-serverless-cache",
            "Description": "A serverless cache.",
            "Status": "deleting",
            "Engine": "redis",
            "MajorEngineVersion": "7",
            "FullEngineVersion": "7.0",
            "SubnetIds": [
                "subnet-xxx8c982",
                "subnet-xxx382f3",
                "subnet-xxxb3e7c0"
            ],
            "SecurityGroupIds": [
                "sg-xxx0c9af"
            ],
            "Endpoint": {
                "Address": "my-serverless-cache-
xxxxx.serverless.use1qa.cache.amazonaws.com",
                "Port": 6379
            },
            "ARN": "arn:aws:elasticache:us-east-1:222222222222:serverlesscache:my-
serverless-cache",
            "SnapshotRetentionLimit": 10,
            "DailySnapshotTime": "11:00"
        }
    }
}
```

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

DeleteServerlessCacheSnapshot

Deletes an existing serverless cache snapshot. Available for Valkey, Redis OSS and Serverless Memcached only.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ServerlessCacheSnapshotName

Idenfitier of the snapshot to be deleted. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: Yes

Response Elements

The following element is returned by the service.

ServerlessCacheSnapshot

The snapshot to be deleted. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: ServerlessCacheSnapshot object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidServerlessCacheSnapshotStateFault

The state of the serverless cache snapshot was not received. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 400

ServerlessCacheSnapshotNotFoundFault

This serverless cache snapshot could not be found or does not exist. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 404

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

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

DeleteSnapshot

Deletes an existing snapshot. When you receive a successful response from this operation, ElastiCache immediately begins deleting the snapshot; you cannot cancel or revert this operation.

> **Note**
>
> This operation is valid for Valkey or Redis OSS only.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

SnapshotName

The name of the snapshot to be deleted.

Type: String

Required: Yes

Response Elements

The following element is returned by the service.

Snapshot

Represents a copy of an entire Valkey or Redis OSS cluster as of the time when the snapshot was taken.

Type: Snapshot object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidSnapshotState

The current state of the snapshot does not allow the requested operation to occur.

HTTP Status Code: 400

SnapshotNotFoundFault

The requested snapshot name does not refer to an existing snapshot.

HTTP Status Code: 404

Examples

DeleteSnapshot

This example illustrates one usage of DeleteSnapshot.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DeleteSnapshot
   &SnapshotName=my-manual-snapshot
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DeleteSnapshotResponse xmlns="http://elasticache.amazonaws.com/doc/2015-02-02/">
   <DeleteSnapshotResult>
      <Snapshot>
         <CacheClusterId>my-redis-primary</CacheClusterId>
```

```
          <Port>6379</Port>
          <CacheNodeType>cache.m1.small</CacheNodeType>
          <CacheParameterGroupName>default.redis2.8</CacheParameterGroupName>
          <Engine>redis</Engine>
          <PreferredAvailabilityZone>us-west-2c</PreferredAvailabilityZone>
          <CacheClusterCreateTime>2015-02-02T18:46:57.972Z</CacheClusterCreateTime>
          <EngineVersion>2.8.6</EngineVersion>
          <SnapshotSource>manual</SnapshotSource>
          <AutoMinorVersionUpgrade>true</AutoMinorVersionUpgrade>
          <PreferredMaintenanceWindow>wed:09:00-wed:10:00</PreferredMaintenanceWindow>
          <SnapshotName>my-manual-snapshot</SnapshotName>
          <SnapshotRetentionLimit>5</SnapshotRetentionLimit>
          <NodeSnapshots>
             <NodeSnapshot>
                <SnapshotCreateTime>2015-02-02T18:54:12Z</SnapshotCreateTime>
                <CacheNodeCreateTime>2015-02-02T18:46:57.972Z</CacheNodeCreateTime>
                <CacheNodeId>0001</CacheNodeId>
                <CacheSize>3 MB</CacheSize>
             </NodeSnapshot>
          </NodeSnapshots>
          <SnapshotStatus>deleting</SnapshotStatus>
          <NumCacheNodes>1</NumCacheNodes>
          <SnapshotWindow>07:30-08:30</SnapshotWindow>
       </Snapshot>
    </DeleteSnapshotResult>
    <ResponseMetadata>
       <RequestId>694d7017-b9d2-11e3-8a16-7978bb24ffdf</RequestId>
    </ResponseMetadata>
</DeleteSnapshotResponse>
```

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

DeleteUser

For Valkey engine version 7.2 onwards and Redis OSS 6.0 onwards: Deletes a user. The user will be removed from all user groups and in turn removed from all replication groups. For more information, see Using Role Based Access Control (RBAC).

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

UserId

The ID of the user.

Type: String

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: Yes

Response Elements

The following elements are returned by the service.

AccessString

Access permissions string used for this user.

Type: String

ARN

The Amazon Resource Name (ARN) of the user.

Type: String

Authentication

Denotes whether the user requires a password to authenticate.

Type: Authentication object

Engine

The options are valkey or redis.

Type: String

Pattern: `[a-zA-Z]*`

MinimumEngineVersion

The minimum engine version required, which is Redis OSS 6.0

Type: String

Status

Indicates the user status. Can be "active", "modifying" or "deleting".

Type: String

UserGroupIds.member.N

Returns a list of the user group IDs the user belongs to.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

UserId

The ID of the user.

Type: String

UserName

The username of the user.

Type: String

Errors

For information about the errors that are common to all actions, see Common Errors.

DefaultUserAssociatedToUserGroup

The default user assigned to the user group.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidUserState

The user is not in active state.

HTTP Status Code: 400

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

UserNotFound

The user does not exist or could not be found.

HTTP Status Code: 404

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

DeleteUserGroup

For Valkey engine version 7.2 onwards and Redis OSS 6.0 onwards: Deletes a user group. The user group must first be disassociated from the replication group before it can be deleted. For more information, see Using Role Based Access Control (RBAC).

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

UserGroupId

The ID of the user group.

Type: String

Required: Yes

Response Elements

The following elements are returned by the service.

ARN

The Amazon Resource Name (ARN) of the user group.

Type: String

Engine

The options are valkey or redis.

Type: String

Pattern: `[a-zA-Z]*`

MinimumEngineVersion

The minimum engine version required, which is Redis OSS 6.0

Type: String

PendingChanges

A list of updates being applied to the user group.

Type: UserGroupPendingChanges object

ReplicationGroups.member.N

A list of replication groups that the user group can access.

Type: Array of strings

ServerlessCaches.member.N

Indicates which serverless caches the specified user group is associated with. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: Array of strings

Status

Indicates user group status. Can be "creating", "active", "modifying", "deleting".

Type: String

UserGroupId

The ID of the user group.

Type: String

UserIds.member.N

The list of user IDs that belong to the user group.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidUserGroupState

The user group is not in an active state.

HTTP Status Code: 400

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

UserGroupNotFound

The user group was not found or does not exist

HTTP Status Code: 404

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

DescribeCacheClusters

Returns information about all provisioned clusters if no cluster identifier is specified, or about a specific cache cluster if a cluster identifier is supplied.

By default, abbreviated information about the clusters is returned. You can use the optional *ShowCacheNodeInfo* flag to retrieve detailed information about the cache nodes associated with the clusters. These details include the DNS address and port for the cache node endpoint.

If the cluster is in the *creating* state, only cluster-level information is displayed until all of the nodes are successfully provisioned.

If the cluster is in the *deleting* state, only cluster-level information is displayed.

If cache nodes are currently being added to the cluster, node endpoint information and creation time for the additional nodes are not displayed until they are completely provisioned. When the cluster state is *available*, the cluster is ready for use.

If cache nodes are currently being removed from the cluster, no endpoint information for the removed nodes is displayed.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheClusterId

The user-supplied cluster identifier. If this parameter is specified, only information about that specific cluster is returned. This parameter isn't case sensitive.

Type: String

Required: No

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified `MaxRecords` value, a marker is included in the response so that the remaining results can be retrieved.

Default: 100

Constraints: minimum 20; maximum 100.

Type: Integer

Required: No

ShowCacheClustersNotInReplicationGroups

An optional flag that can be included in the `DescribeCacheCluster` request to show only nodes (API/CLI: clusters) that are not members of a replication group. In practice, this means Memcached and single node Valkey or Redis OSS clusters.

Type: Boolean

Required: No

ShowCacheNodeInfo

An optional flag that can be included in the `DescribeCacheCluster` request to retrieve information about the individual cache nodes.

Type: Boolean

Required: No

Response Elements

The following elements are returned by the service.

CacheClusters.CacheCluster.N

A list of clusters. Each item in the list contains detailed information about one cluster.

Type: Array of CacheCluster objects

Marker

Provides an identifier to allow retrieval of paginated results.

Type: String

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterNotFound

The requested cluster ID does not refer to an existing cluster.

HTTP Status Code: 404

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

Examples

DescribeCacheClusters

This example illustrates one usage of DescribeCacheClusters.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DescribeCacheClusters
   &MaxRecords=100
   &ShowCacheNodeInfo=false
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
```

```
    &Timestamp=20150202T192317Z
    &X-Amz-Credential=<credential>
```

Sample Response

```
<DescribeCacheClustersResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
  <DescribeCacheClustersResult>
    <CacheClusters>
      <CacheCluster>
        <CacheParameterGroup>
          <ParameterApplyStatus>in-sync</ParameterApplyStatus>
          <CacheParameterGroupName>default.memcached1.4</CacheParameterGroupName>
          <CacheNodeIdsToReboot/>
        </CacheParameterGroup>
        <CacheClusterId>simcoprod42</CacheClusterId>
        <CacheClusterStatus>available</CacheClusterStatus>
        <ConfigurationEndpoint>
          <Port>11211</Port>
          <Address>simcoprod42.m2st2p.cfg.cache.amazonaws.com</Address>
        </ConfigurationEndpoint>
        <ClientDownloadLandingPage>
          https://console.aws.amazon.com/elasticache/home#client-download:
        </ClientDownloadLandingPage>
        <CacheNodeType>cache.m1.large</CacheNodeType>
        <Engine>memcached</Engine>
        <PendingModifiedValues/>
        <PreferredAvailabilityZone>us-west-2c</PreferredAvailabilityZone>
        <CacheClusterCreateTime>2015-02-02T01:21:46.607Z</CacheClusterCreateTime>
        <EngineVersion>1.4.5</EngineVersion>
        <AutoMinorVersionUpgrade>true</AutoMinorVersionUpgrade>
        <PreferredMaintenanceWindow>fri:08:30-fri:09:30</PreferredMaintenanceWindow>
        <CacheSecurityGroups>
          <CacheSecurityGroup>
            <CacheSecurityGroupName>default</CacheSecurityGroupName>
            <Status>active</Status>
          </CacheSecurityGroup>
        </CacheSecurityGroups>
        <NotificationConfiguration>
          <TopicStatus>active</TopicStatus>
          <TopicArn>arn:aws:sns:us-west-2:123456789012:ElastiCacheNotifications</
TopicArn>
```

```
        </NotificationConfiguration>
        <NumCacheNodes>6</NumCacheNodes>
      </CacheCluster>
    </CacheClusters>
  </DescribeCacheClustersResult>
  <ResponseMetadata>
    <RequestId>f270d58f-b7fb-11e0-9326-b7275b9d4a6c</RequestId>
  </ResponseMetadata>
</DescribeCacheClustersResponse>
```

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

DescribeCacheEngineVersions

Returns a list of the available cache engines and their versions.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheParameterGroupFamily

The name of a specific cache parameter group family to return details for.

Valid values are: `memcached1.4` | `memcached1.5` | `memcached1.6` | `redis2.6` | `redis2.8` | `redis3.2` | `redis4.0` | `redis5.0` | `redis6.x` | `redis6.2` | `redis7` | `valkey7`

Constraints:

- Must be 1 to 255 alphanumeric characters
- First character must be a letter
- Cannot end with a hyphen or contain two consecutive hyphens

Type: String

Required: No

DefaultOnly

If `true`, specifies that only the default version of the specified engine or engine and major version combination is to be returned.

Type: Boolean

Required: No

Engine

The cache engine to return. Valid values: `memcached` | `redis`

Type: String

Required: No

EngineVersion

The cache engine version to return.

Example: `1.4.14`

Type: String

Required: No

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified `MaxRecords` value, a marker is included in the response so that the remaining results can be retrieved.

Default: 100

Constraints: minimum 20; maximum 100.

Type: Integer

Required: No

Response Elements

The following elements are returned by the service.

CacheEngineVersions.CacheEngineVersion.N

A list of cache engine version details. Each element in the list contains detailed information about one cache engine version.

Type: Array of CacheEngineVersion objects

Marker

Provides an identifier to allow retrieval of paginated results.

Type: String

Errors

For information about the errors that are common to all actions, see Common Errors.

Examples

DescribeCacheEngineVersions

This example illustrates one usage of DescribeCacheEngineVersions.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DescribeCacheEngineVersions
   &MaxRecords=100
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DescribeCacheEngineVersionsResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
  <DescribeCacheEngineVersionsResult>
    <CacheEngineVersions>
      <CacheEngineVersion>
        <CacheParameterGroupFamily>memcached1.4</CacheParameterGroupFamily>
        <Engine>memcached</Engine>
        <CacheEngineVersionDescription>memcached version 1.4.14</
CacheEngineVersionDescription>
        <CacheEngineDescription>memcached</CacheEngineDescription>
        <EngineVersion>1.4.14</EngineVersion>
      </CacheEngineVersion>
      <CacheEngineVersion>
        <CacheParameterGroupFamily>memcached1.4</CacheParameterGroupFamily>
        <Engine>memcached</Engine>
```

```
        <CacheEngineVersionDescription>memcached version 1.4.5</
CacheEngineVersionDescription>
        <CacheEngineDescription>memcached</CacheEngineDescription>
        <EngineVersion>1.4.5</EngineVersion>
      </CacheEngineVersion>
    </CacheEngineVersions>
  </DescribeCacheEngineVersionsResult>
  <ResponseMetadata>
    <RequestId>a6ac9ad2-f8a4-11e1-a4d1-a345e5370093</RequestId>
  </ResponseMetadata>
</DescribeCacheEngineVersionsResponse>
```

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

DescribeCacheParameterGroups

Returns a list of cache parameter group descriptions. If a cache parameter group name is specified, the list contains only the descriptions for that group.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheParameterGroupName

The name of a specific cache parameter group to return details for.

Type: String

Required: No

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified `MaxRecords` value, a marker is included in the response so that the remaining results can be retrieved.

Default: 100

Constraints: minimum 20; maximum 100.

Type: Integer

Required: No

Response Elements

The following elements are returned by the service.

CacheParameterGroups.CacheParameterGroup.N

A list of cache parameter groups. Each element in the list contains detailed information about one cache parameter group.

Type: Array of CacheParameterGroup objects

Marker

Provides an identifier to allow retrieval of paginated results.

Type: String

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheParameterGroupNotFound

The requested cache parameter group name does not refer to an existing cache parameter group.

HTTP Status Code: 404

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

Examples

DescribeCacheParameterGroups

This example illustrates one usage of DescribeCacheParameterGroups.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DescribeCacheParameterGroups
   &MaxRecords=100
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DescribeCacheParameterGroupsResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
  <DescribeCacheParameterGroupsResult>
    <CacheParameterGroups>
      <CacheParameterGroup>
        <CacheParameterGroupName>default.memcached1.4</CacheParameterGroupName>
        <CacheParameterGroupFamily>memcached1.4</CacheParameterGroupFamily>
        <Description>Default parameter group for memcached1.4</Description>
      </CacheParameterGroup>
      <CacheParameterGroup>
        <CacheParameterGroupName>mycacheparametergroup</CacheParameterGroupName>
        <CacheParameterGroupFamily>memcached1.4</CacheParameterGroupFamily>
        <Description>My cache parameter group</Description>
      </CacheParameterGroup>
      <CacheParameterGroup>
        <CacheParameterGroupName>mycacheparametergroup1</CacheParameterGroupName>
        <CacheParameterGroupFamily>memcached1.4</CacheParameterGroupFamily>
        <Description>My first cache parameter group</Description>
      </CacheParameterGroup>
      <CacheParameterGroup>
        <CacheParameterGroupName>mycacheparametergroup3</CacheParameterGroupName>
        <CacheParameterGroupFamily>memcached1.4</CacheParameterGroupFamily>
        <Description>My first cache parameter group</Description>
      </CacheParameterGroup>
    </CacheParameterGroups>
  </DescribeCacheParameterGroupsResult>
  <ResponseMetadata>
    <RequestId>7193fbb8-b7fc-11e0-9b0b-a9261be2b354</RequestId>
  </ResponseMetadata>
</DescribeCacheParameterGroupsResponse>
```

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

DescribeCacheParameters

Returns the detailed parameter list for a particular cache parameter group.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheParameterGroupName

The name of a specific cache parameter group to return details for.

Type: String

Required: Yes

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified `MaxRecords` value, a marker is included in the response so that the remaining results can be retrieved.

Default: 100

Constraints: minimum 20; maximum 100.

Type: Integer

Required: No

Source

The parameter types to return.

Valid values: `user` | `system` | `engine-default`

Type: String

Required: No

Response Elements

The following elements are returned by the service.

CacheNodeTypeSpecificParameters.CacheNodeTypeSpecificParameter.N

A list of parameters specific to a particular cache node type. Each element in the list contains detailed information about one parameter.

Type: Array of CacheNodeTypeSpecificParameter objects

Marker

Provides an identifier to allow retrieval of paginated results.

Type: String

Parameters.Parameter.N

A list of Parameter instances.

Type: Array of Parameter objects

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheParameterGroupNotFound

The requested cache parameter group name does not refer to an existing cache parameter group.

HTTP Status Code: 404

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

Examples

DescribeCacheParameters

Some of the output has been omitted for brevity.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DescribeCacheParameters
   &CacheParameterGroupName=default.memcached1.4
   &MaxRecords=100
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DescribeCacheParametersResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
  <DescribeCacheParametersResult>
    <CacheNodeTypeSpecificParameters>
      <CacheNodeTypeSpecificParameter>
        <CacheNodeTypeSpecificValues>
          <CacheNodeTypeSpecificValue>
            <CacheNodeType>cache.c1.xlarge</CacheNodeType>
            <Value>6000</Value>
          </CacheNodeTypeSpecificValue>

 (...output omitted...)

        </CacheNodeTypeSpecificValues>
        <DataType>integer</DataType>
        <Source>system</Source>
```

```
        <IsModifiable>false</IsModifiable>
        <Description>The maximum configurable amount of memory to use to store items,
 in megabytes.</Description>
        <AllowedValues>1-100000</AllowedValues>
        <ParameterName>max_cache_memory</ParameterName>
        <MinimumEngineVersion>1.4.5</MinimumEngineVersion>
      </CacheNodeTypeSpecificParameter>
      <CacheNodeTypeSpecificParameter>

 (...output omitted...)

      </CacheNodeTypeSpecificParameter>
    </CacheNodeTypeSpecificParameters>
    <Parameters>
      <Parameter>
        <ParameterValue>1024</ParameterValue>
        <DataType>integer</DataType>
        <Source>system</Source>
        <IsModifiable>false</IsModifiable>
        <Description>The backlog queue limit.</Description>
        <AllowedValues>1-10000</AllowedValues>
        <ParameterName>backlog_queue_limit</ParameterName>
        <MinimumEngineVersion>1.4.5</MinimumEngineVersion>
      </Parameter>

 (...output omitted...)

    </Parameters>
  </DescribeCacheParametersResult>
  <ResponseMetadata>
    <RequestId>0c507368-b7fe-11e0-9326-b7275b9d4a6c</RequestId>
  </ResponseMetadata>
</DescribeCacheParametersResponse>
```

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

DescribeCacheSecurityGroups

Returns a list of cache security group descriptions. If a cache security group name is specified, the list contains only the description of that group. This applicable only when you have ElastiCache in Classic setup

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheSecurityGroupName

The name of the cache security group to return details for.

Type: String

Required: No

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified `MaxRecords` value, a marker is included in the response so that the remaining results can be retrieved.

Default: 100

Constraints: minimum 20; maximum 100.

Type: Integer

Required: No

Response Elements

The following elements are returned by the service.

CacheSecurityGroups.CacheSecurityGroup.N

A list of cache security groups. Each element in the list contains detailed information about one group.

Type: Array of CacheSecurityGroup objects

Marker

Provides an identifier to allow retrieval of paginated results.

Type: String

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheSecurityGroupNotFound

The requested cache security group name does not refer to an existing cache security group.

HTTP Status Code: 404

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

Examples

DescribeCacheSecurityGroups

This example illustrates one usage of DescribeCacheSecurityGroups.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DescribeCacheSecurityGroups
   &MaxRecords=100
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DescribeCacheSecurityGroupsResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
  <DescribeCacheSecurityGroupsResult>
    <CacheSecurityGroups>
      <CacheSecurityGroup>
        <EC2SecurityGroups/>
        <CacheSecurityGroupName>default</CacheSecurityGroupName>
        <OwnerId>123456789012</OwnerId>
        <Description>default</Description>
      </CacheSecurityGroup>
      <CacheSecurityGroup>
        <EC2SecurityGroups/>
        <CacheSecurityGroupName>mycachesecuritygroup</CacheSecurityGroupName>
        <OwnerId>123456789012</OwnerId>
        <Description>My Security Group</Description>
      </CacheSecurityGroup>
    </CacheSecurityGroups>
  </DescribeCacheSecurityGroupsResult>
  <ResponseMetadata>
    <RequestId>a95360ae-b7fc-11e0-9326-b7275b9d4a6c</RequestId>
  </ResponseMetadata>
</DescribeCacheSecurityGroupsResponse>
```

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

DescribeCacheSubnetGroups

Returns a list of cache subnet group descriptions. If a subnet group name is specified, the list contains only the description of that group. This is applicable only when you have ElastiCache in VPC setup. All ElastiCache clusters now launch in VPC by default.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheSubnetGroupName

The name of the cache subnet group to return details for.

Type: String

Required: No

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified `MaxRecords` value, a marker is included in the response so that the remaining results can be retrieved.

Default: 100

Constraints: minimum 20; maximum 100.

Type: Integer

Required: No

Response Elements

The following elements are returned by the service.

CacheSubnetGroups.CacheSubnetGroup.N

A list of cache subnet groups. Each element in the list contains detailed information about one group.

Type: Array of CacheSubnetGroup objects

Marker

Provides an identifier to allow retrieval of paginated results.

Type: String

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheSubnetGroupNotFoundFault

The requested cache subnet group name does not refer to an existing cache subnet group.

HTTP Status Code: 400

Examples

DescribeCacheSubnetGroups

Some of the output has been omitted for brevity.

Sample Request

```
https://elasticache.amazonaws.com/
   ?Action=DescribeCacheSubnetGroups
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
```

```
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DescribeCacheSubnetGroupsResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
        <DescribeCacheSubnetGroupsResult>
            <CacheSubnetGroups>
                <CacheSubnetGroup>
                    <VpcId>990524496922</VpcId>
                    <CacheSubnetGroupDescription>description</
CacheSubnetGroupDescription>
                    <CacheSubnetGroupName>subnet_grp1</CacheSubnetGroupName>
                    <Subnets>
                        <Subnet>
                            <SubnetStatus>Active</SubnetStatus>
                            <SubnetIdentifier>subnet-7c5b4115</SubnetIdentifier>
                            <SubnetAvailabilityZone>
                                <Name>us-west-2c</Name>
                            </SubnetAvailabilityZone>
                        </Subnet>
                        <Subnet>
                            <SubnetStatus>Active</SubnetStatus>
                            <SubnetIdentifier>subnet-7b5b4112</SubnetIdentifier>
                            <SubnetAvailabilityZone>
                                <Name>us-west-2b</Name>
                            </SubnetAvailabilityZone>
                        </Subnet>
                        <Subnet>
                            <SubnetStatus>Active</SubnetStatus>
                            <SubnetIdentifier>subnet-3ea6bd57</SubnetIdentifier>
                            <SubnetAvailabilityZone>
                                <Name>us-west-2c</Name>
                            </SubnetAvailabilityZone>
                        </Subnet>
                    </Subnets>
                </CacheSubnetGroup>

 (...output omitted...)

            </CacheSubnetGroups>
        </DescribeCacheSubnetGroupsResult>
        <ResponseMetadata>
```

```
            <RequestId>31d0faee-229b-11e1-81f1-df3a2a803dad</RequestId>
        </ResponseMetadata>
    </DescribeCacheSubnetGroupsResponse>
```

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

DescribeEngineDefaultParameters

Returns the default engine and system parameter information for the specified cache engine.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheParameterGroupFamily

The name of the cache parameter group family.

Valid values are: `memcached1.4` | `memcached1.5` | `memcached1.6` | `redis2.6` | `redis2.8` | `redis3.2` | `redis4.0` | `redis5.0` | `redis6.x` | `redis6.2` | `redis7`

Type: String

Required: Yes

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified `MaxRecords` value, a marker is included in the response so that the remaining results can be retrieved.

Default: 100

Constraints: minimum 20; maximum 100.

Type: Integer

Required: No

Response Elements

The following element is returned by the service.

EngineDefaults

Represents the output of a `DescribeEngineDefaultParameters` operation.

Type: EngineDefaults object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

Examples

DescribeEngineDefaultParameters

Some of the output has been omitted for brevity.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DescribeEngineDefaultParameters
   &CacheParameterGroupFamily=memcached1.4
   &MaxRecords=100
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
```

```
  &X-Amz-Credential=<credential>
```

Sample Response

```
<DescribeEngineDefaultParametersResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
  <DescribeEngineDefaultParametersResult>
    <EngineDefaults>
      <CacheParameterGroupFamily>memcached1.4</CacheParameterGroupFamily>
      <Parameters>
        <Parameter>
          <ParameterValue>1024</ParameterValue>
          <DataType>integer</DataType>
          <Source>system</Source>
          <IsModifiable>false</IsModifiable>
          <Description>The backlog queue limit.</Description>
          <AllowedValues>1-10000</AllowedValues>
          <ParameterName>backlog_queue_limit</ParameterName>
          <MinimumEngineVersion>1.4.5</MinimumEngineVersion>
        </Parameter>
        <Parameter>

 (...output omitted...)

        </Parameter>
      </Parameters>
      <CacheNodeTypeSpecificParameters>
        <CacheNodeTypeSpecificParameter>
          <CacheNodeTypeSpecificValues>
            <CacheNodeTypeSpecificValue>
              <CacheNodeType>cache.c1.xlarge</CacheNodeType>
              <Value>6000</Value>
            </CacheNodeTypeSpecificValue>

 (...output omitted...)

          </CacheNodeTypeSpecificValues>
          <DataType>integer</DataType>
          <Source>system</Source>
          <IsModifiable>false</IsModifiable>
          <Description>The maximum configurable amount of memory to use to store items,
 in megabytes.</Description>
```

```
          <AllowedValues>1-100000</AllowedValues>
          <ParameterName>max_cache_memory</ParameterName>
          <MinimumEngineVersion>1.4.5</MinimumEngineVersion>
        </CacheNodeTypeSpecificParameter>

 (...output omitted...)

      </CacheNodeTypeSpecificParameters>
    </EngineDefaults>
  </DescribeEngineDefaultParametersResult>
  <ResponseMetadata>
    <RequestId>061282fe-b7fd-11e0-9326-b7275b9d4a6c</RequestId>
  </ResponseMetadata>
</DescribeEngineDefaultParametersResponse>
```

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

DescribeEvents

Returns events related to clusters, cache security groups, and cache parameter groups. You can obtain events specific to a particular cluster, cache security group, or cache parameter group by providing the name as a parameter.

By default, only the events occurring within the last hour are returned; however, you can retrieve up to 14 days' worth of events if necessary.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

Duration

The number of minutes worth of events to retrieve.

Type: Integer

Required: No

EndTime

The end of the time interval for which to retrieve events, specified in ISO 8601 format.

Example: 2017-03-30T07:03:49.555Z

Type: Timestamp

Required: No

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified `MaxRecords` value, a marker is included in the response so that the remaining results can be retrieved.

Default: 100

Constraints: minimum 20; maximum 100.

Type: Integer

Required: No

SourceIdentifier

The identifier of the event source for which events are returned. If not specified, all sources are included in the response.

Type: String

Required: No

SourceType

The event source to retrieve events for. If no value is specified, all events are returned.

Type: String

Valid Values: `cache-cluster | cache-parameter-group | cache-security-group | cache-subnet-group | replication-group | serverless-cache | serverless-cache-snapshot | user | user-group`

Required: No

StartTime

The beginning of the time interval to retrieve events for, specified in ISO 8601 format.

Example: 2017-03-30T07:03:49.555Z

Type: Timestamp

Required: No

Response Elements

The following elements are returned by the service.

Events.Event.N

A list of events. Each element in the list contains detailed information about one event.

Type: Array of Event objects

Marker

Provides an identifier to allow retrieval of paginated results.

Type: String

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

Examples

DescribeEvents

Some of the output has been omitted for brevity.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DescribeEvents
```

```
   &MaxRecords=100
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DescribeEventsResponse xmlns="http://elasticache.amazonaws.com/doc/2015-02-02/">
    <DescribeEventsResult>
        <Events>
            <Event>
                <Message>Cache cluster created</Message>
                <SourceType>cache-cluster</SourceType>
                <Date>2015-02-02T18:22:18.202Z</Date>
                <SourceIdentifier>my-redis-primary</SourceIdentifier>
            </Event>

 (...output omitted...)

        </Events>
    </DescribeEventsResult>
    <ResponseMetadata>
        <RequestId>e21c81b4-b9cd-11e3-8a16-7978bb24ffdf</RequestId>
    </ResponseMetadata>
</DescribeEventsResponse>
```

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

DescribeGlobalReplicationGroups

Returns information about a particular global replication group. If no identifier is specified, returns information about all Global datastores.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

GlobalReplicationGroupId

The name of the Global datastore

Type: String

Required: No

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified MaxRecords value, a marker is included in the response so that the remaining results can be retrieved.

Type: Integer

Required: No

ShowMemberInfo

Returns the list of members that comprise the Global datastore.

Type: Boolean

Required: No

Response Elements

The following elements are returned by the service.

GlobalReplicationGroups.GlobalReplicationGroup.N

Indicates the slot configuration and global identifier for each slice group.

Type: Array of GlobalReplicationGroup objects

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by MaxRecords. >

Type: String

Errors

For information about the errors that are common to all actions, see Common Errors.

GlobalReplicationGroupNotFoundFault

The Global datastore does not exist

HTTP Status Code: 404

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

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

DescribeReplicationGroups

Returns information about a particular replication group. If no identifier is specified, `DescribeReplicationGroups` returns information about all replication groups.

> **Note**
>
> This operation is valid for Valkey or Redis OSS only.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified `MaxRecords` value, a marker is included in the response so that the remaining results can be retrieved.

Default: 100

Constraints: minimum 20; maximum 100.

Type: Integer

Required: No

ReplicationGroupId

The identifier for the replication group to be described. This parameter is not case sensitive.

If you do not specify this parameter, information about all replication groups is returned.

Type: String

Required: No

Response Elements

The following elements are returned by the service.

Marker

Provides an identifier to allow retrieval of paginated results.

Type: String

ReplicationGroups.ReplicationGroup.N

A list of replication groups. Each item in the list contains detailed information about one replication group.

Type: Array of ReplicationGroup objects

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

Examples

DescribeReplicationGroups

This example illustrates one usage of DescribeReplicationGroups.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DescribeReplicationGroups
   &MaxRecords=100
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DescribeReplicationGroupsResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
   <DescribeReplicationGroupsResult>
      <ReplicationGroups>
         <ReplicationGroup>
            <SnapshottingClusterId>my-redis-primary</SnapshottingClusterId>
            <MemberClusters>
               <ClusterId>my-redis-primary</ClusterId>
            </MemberClusters>
            <NodeGroups>
               <NodeGroup>
                  <NodeGroupId>0001</NodeGroupId>
                  <PrimaryEndpoint>
                     <Port>6379</Port>
                     <Address>my-repgroup.q68zge.ng.0001.use1devo.elmo-
dev.amazonaws.com</Address>
                  </PrimaryEndpoint>
                  <Status>available</Status>
                  <NodeGroupMembers>
                     <NodeGroupMember>
                        <CacheClusterId>my-redis-primary</CacheClusterId>
```

```
                        <ReadEndpoint>
                           <Port>6379</Port>
                           <Address>my-redis-primary.q68zge.0001.use1devo.elmo-
dev.amazonaws.com</Address>
                        </ReadEndpoint>
                        <PreferredAvailabilityZone>us-west-2c</
PreferredAvailabilityZone>
                        <CacheNodeId>0001</CacheNodeId>
                        <CurrentRole>primary</CurrentRole>
                     </NodeGroupMember>
                  </NodeGroupMembers>
               </NodeGroup>
            </NodeGroups>
            <ReplicationGroupId>my-repgroup</ReplicationGroupId>
            <Status>available</Status>
            <PendingModifiedValues />
            <Description>My replication group</Description>
         </ReplicationGroup>
      </ReplicationGroups>
   </DescribeReplicationGroupsResult>
   <ResponseMetadata>
      <RequestId>144745b0-b9d3-11e3-8a16-7978bb24ffdf</RequestId>
   </ResponseMetadata>
</DescribeReplicationGroupsResponse>
```

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

DescribeReservedCacheNodes

Returns information about reserved cache nodes for this account, or about a specified reserved cache node.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheNodeType

The cache node type filter value. Use this parameter to show only those reservations matching the specified cache node type.

The following node types are supported by ElastiCache. Generally speaking, the current generation types provide more memory and computational power at lower cost when compared to their equivalent previous generation counterparts.

- General purpose:
  - Current generation:

    M7g node types: `cache.m7g.large`, `cache.m7g.xlarge`, `cache.m7g.2xlarge`, `cache.m7g.4xlarge`, `cache.m7g.8xlarge`, `cache.m7g.12xlarge`, `cache.m7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    M6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.m6g.large`, `cache.m6g.xlarge`, `cache.m6g.2xlarge`, `cache.m6g.4xlarge`, `cache.m6g.8xlarge`, `cache.m6g.12xlarge`, `cache.m6g.16xlarge`

    M5 node types: `cache.m5.large`, `cache.m5.xlarge`, `cache.m5.2xlarge`, `cache.m5.4xlarge`, `cache.m5.12xlarge`, `cache.m5.24xlarge`

    M4 node types: `cache.m4.large`, `cache.m4.xlarge`, `cache.m4.2xlarge`, `cache.m4.4xlarge`, `cache.m4.10xlarge`

T4g node types (available only for Redis OSS engine version 5.0.6 onward and Memcached engine version 1.5.16 onward): `cache.t4g.micro`, `cache.t4g.small`, `cache.t4g.medium`

T3 node types: `cache.t3.micro`, `cache.t3.small`, `cache.t3.medium`

T2 node types: `cache.t2.micro`, `cache.t2.small`, `cache.t2.medium`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  T1 node types: `cache.t1.micro`

  M1 node types: `cache.m1.small`, `cache.m1.medium`, `cache.m1.large`, `cache.m1.xlarge`

  M3 node types: `cache.m3.medium`, `cache.m3.large`, `cache.m3.xlarge`, `cache.m3.2xlarge`

- Compute optimized:
  - Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

    C1 node types: `cache.c1.xlarge`

- Memory optimized:
  - Current generation:

    R7g node types: `cache.r7g.large`, `cache.r7g.xlarge`, `cache.r7g.2xlarge`, `cache.r7g.4xlarge`, `cache.r7g.8xlarge`, `cache.r7g.12xlarge`, `cache.r7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    R6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.r6g.large`, `cache.r6g.xlarge`,

`cache.r6g.2xlarge`, `cache.r6g.4xlarge`, `cache.r6g.8xlarge`, `cache.r6g.12xlarge`, `cache.r6g.16xlarge`

R5 node types: `cache.r5.large`, `cache.r5.xlarge`, `cache.r5.2xlarge`, `cache.r5.4xlarge`, `cache.r5.12xlarge`, `cache.r5.24xlarge`

R4 node types: `cache.r4.large`, `cache.r4.xlarge`, `cache.r4.2xlarge`, `cache.r4.4xlarge`, `cache.r4.8xlarge`, `cache.r4.16xlarge`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  M2 node types: `cache.m2.xlarge`, `cache.m2.2xlarge`, `cache.m2.4xlarge`

  R3 node types: `cache.r3.large`, `cache.r3.xlarge`, `cache.r3.2xlarge`, `cache.r3.4xlarge`, `cache.r3.8xlarge`

Additional node type info

- All current generation instance types are created in Amazon VPC by default.
- Valkey or Redis OSS append-only files (AOF) are not supported for T1 or T2 instances.
- Valkey or Redis OSS Multi-AZ with automatic failover is not supported on T1 instances.
- The configuration variables `appendonly` and `appendfsync` are not supported on Valkey, or on Redis OSS version 2.8.22 and later.

Type: String

Required: No

Duration

The duration filter value, specified in years or seconds. Use this parameter to show only reservations for this duration.

Valid Values: `1 | 3 | 31536000 | 94608000`

Type: String

Required: No

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified `MaxRecords` value, a marker is included in the response so that the remaining results can be retrieved.

Default: 100

Constraints: minimum 20; maximum 100.

Type: Integer

Required: No

OfferingType

The offering type filter value. Use this parameter to show only the available offerings matching the specified offering type.

Valid values: `"Light Utilization"|"Medium Utilization"|"Heavy Utilization"|"All Upfront"|"Partial Upfront"| "No Upfront"`

Type: String

Required: No

ProductDescription

The product description filter value. Use this parameter to show only those reservations matching the specified product description.

Type: String

Required: No

ReservedCacheNodeId

The reserved cache node identifier filter value. Use this parameter to show only the reservation that matches the specified reservation ID.

Type: String

Required: No

ReservedCacheNodesOfferingId

The offering identifier filter value. Use this parameter to show only purchased reservations matching the specified offering identifier.

Type: String

Required: No

Response Elements

The following elements are returned by the service.

Marker

Provides an identifier to allow retrieval of paginated results.

Type: String

ReservedCacheNodes.ReservedCacheNode.N

A list of reserved cache nodes. Each element in the list contains detailed information about one node.

Type: Array of ReservedCacheNode objects

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

ReservedCacheNodeNotFound

The requested reserved cache node was not found.

HTTP Status Code: 404

Examples

DescribeReservedCacheNodes

This example illustrates one usage of DescribeReservedCacheNodes.

Sample Request

```
https://elasticache.amazonaws.com/
   ?Action=DescribeReservedCacheNodes
   &ReservedCacheNodeId=customerSpecifiedID
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DescribeReservedCacheNodesResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
  <DescribeReservedCacheNodesResult>
    <ReservedCacheNodes>
      <ReservedCacheNode>
        <OfferingType>Medium Utilization</OfferingType>
        <RecurringCharges/>
        <ProductDescription>memcached</ProductDescription>
        <ReservedCacheNodesOfferingId>649fd0c8-cf6d-47a0-bfa6-060f8e75e95f</
ReservedCacheNodesOfferingId>
```

```
        <State>payment-failed</State>
        <ReservedCacheNodeId>myreservationid</ReservedCacheNodeId>
        <CacheNodeCount>1</CacheNodeCount>
        <StartTime>2010-12-15T00:25:14.131Z</StartTime>
        <Duration>31536000</Duration>
        <FixedPrice>227.5</FixedPrice>
        <UsagePrice>0.046</UsagePrice>
        <CacheNodeType>cache.m1.small</CacheNodeType>
      </ReservedCacheNode>
  </DescribeReservedCacheNodesResult>
  <ResponseMetadata>
    <RequestId>c695119b-2961-11e1-bd06-6fe008f046c3</RequestId>
  </ResponseMetadata>
</DescribeReservedCacheNodesResponse>
```

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

DescribeReservedCacheNodesOfferings

Lists available reserved cache node offerings.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheNodeType

The cache node type filter value. Use this parameter to show only the available offerings matching the specified cache node type.

The following node types are supported by ElastiCache. Generally speaking, the current generation types provide more memory and computational power at lower cost when compared to their equivalent previous generation counterparts.

- General purpose:
  - Current generation:

    M7g node types: `cache.m7g.large`, `cache.m7g.xlarge`, `cache.m7g.2xlarge`, `cache.m7g.4xlarge`, `cache.m7g.8xlarge`, `cache.m7g.12xlarge`, `cache.m7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    M6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.m6g.large`, `cache.m6g.xlarge`, `cache.m6g.2xlarge`, `cache.m6g.4xlarge`, `cache.m6g.8xlarge`, `cache.m6g.12xlarge`, `cache.m6g.16xlarge`

    M5 node types: `cache.m5.large`, `cache.m5.xlarge`, `cache.m5.2xlarge`, `cache.m5.4xlarge`, `cache.m5.12xlarge`, `cache.m5.24xlarge`

    M4 node types: `cache.m4.large`, `cache.m4.xlarge`, `cache.m4.2xlarge`, `cache.m4.4xlarge`, `cache.m4.10xlarge`

T4g node types (available only for Redis OSS engine version 5.0.6 onward and Memcached engine version 1.5.16 onward): `cache.t4g.micro`, `cache.t4g.small`, `cache.t4g.medium`

T3 node types: `cache.t3.micro`, `cache.t3.small`, `cache.t3.medium`

T2 node types: `cache.t2.micro`, `cache.t2.small`, `cache.t2.medium`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  T1 node types: `cache.t1.micro`

  M1 node types: `cache.m1.small`, `cache.m1.medium`, `cache.m1.large`, `cache.m1.xlarge`

  M3 node types: `cache.m3.medium`, `cache.m3.large`, `cache.m3.xlarge`, `cache.m3.2xlarge`

- Compute optimized:
  - Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

    C1 node types: `cache.c1.xlarge`

- Memory optimized:
  - Current generation:

    R7g node types: `cache.r7g.large`, `cache.r7g.xlarge`, `cache.r7g.2xlarge`, `cache.r7g.4xlarge`, `cache.r7g.8xlarge`, `cache.r7g.12xlarge`, `cache.r7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    R6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.r6g.large`, `cache.r6g.xlarge`,

`cache.r6g.2xlarge`, `cache.r6g.4xlarge`, `cache.r6g.8xlarge`, `cache.r6g.12xlarge`, `cache.r6g.16xlarge`

R5 node types: `cache.r5.large`, `cache.r5.xlarge`, `cache.r5.2xlarge`, `cache.r5.4xlarge`, `cache.r5.12xlarge`, `cache.r5.24xlarge`

R4 node types: `cache.r4.large`, `cache.r4.xlarge`, `cache.r4.2xlarge`, `cache.r4.4xlarge`, `cache.r4.8xlarge`, `cache.r4.16xlarge`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  M2 node types: `cache.m2.xlarge`, `cache.m2.2xlarge`, `cache.m2.4xlarge`

  R3 node types: `cache.r3.large`, `cache.r3.xlarge`, `cache.r3.2xlarge`, `cache.r3.4xlarge`, `cache.r3.8xlarge`

Additional node type info

- All current generation instance types are created in Amazon VPC by default.
- Valkey or Redis OSS append-only files (AOF) are not supported for T1 or T2 instances.
- Valkey or Redis OSS Multi-AZ with automatic failover is not supported on T1 instances.
- The configuration variables `appendonly` and `appendfsync` are not supported on Valkey, or on Redis OSS version 2.8.22 and later.

Type: String

Required: No

Duration

Duration filter value, specified in years or seconds. Use this parameter to show only reservations for a given duration.

Valid Values: `1 | 3 | 31536000 | 94608000`

Type: String

Required: No

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified `MaxRecords` value, a marker is included in the response so that the remaining results can be retrieved.

Default: 100

Constraints: minimum 20; maximum 100.

Type: Integer

Required: No

OfferingType

The offering type filter value. Use this parameter to show only the available offerings matching the specified offering type.

Valid Values: `"Light Utilization"|"Medium Utilization"|"Heavy Utilization" |"All Upfront"|"Partial Upfront"| "No Upfront"`

Type: String

Required: No

ProductDescription

The product description filter value. Use this parameter to show only the available offerings matching the specified product description.

Type: String

Required: No

ReservedCacheNodesOfferingId

The offering identifier filter value. Use this parameter to show only the available offering that matches the specified reservation identifier.

Example: `438012d3-4052-4cc7-b2e3-8d3372e0e706`

Type: String

Required: No

Response Elements

The following elements are returned by the service.

Marker

Provides an identifier to allow retrieval of paginated results.

Type: String

ReservedCacheNodesOfferings.ReservedCacheNodesOffering.N

A list of reserved cache node offerings. Each element in the list contains detailed information about one offering.

Type: Array of ReservedCacheNodesOffering objects

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

ReservedCacheNodesOfferingNotFound

The requested cache node offering does not exist.

HTTP Status Code: 404

Examples

DescribeReservedCacheNodesOfferings

This example illustrates one usage of DescribeReservedCacheNodesOfferings.

Sample Request

```
https://elasticache.amazonaws.com/
   ?Action=DescribeReservedCacheNodesOfferings
   &ReservedCacheNodesOfferingId=438012d3-4052-4cc7-b2e3-8d3372e0e706
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DescribeReservedCacheNodesOfferingsResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
  <DescribeReservedCacheNodesOfferingsResult>
    <ReservedCacheNodesOfferings>
      <ReservedCacheNodesOffering>
        <Duration>31536000</Duration>
        <OfferingType>Heavy Utilization</OfferingType>
        <RecurringCharges>
          <RecurringCharge>
            <RecurringChargeFrequency>Hourly</RecurringChargeFrequency>
            <RecurringChargeAmount>0.123</RecurringChargeAmount>
          </RecurringCharge>
        </RecurringCharges>
        <FixedPrice>162.0</FixedPrice>
        <ProductDescription>memcached</ProductDescription>
        <UsagePrice>0.0</UsagePrice>
        <ReservedCacheNodesOfferingId>SampleOfferingId</ReservedCacheNodesOfferingId>
```

```
        <CacheNodeType>cache.m1.small</CacheNodeType>
      </ReservedCacheNodesOffering>
    </ReservedCacheNodesOfferings>
  </DescribeReservedCacheNodesOfferingsResult>
  <ResponseMetadata>
    <RequestId>521b420a-2961-11e1-bd06-6fe008f046c3</RequestId>
  </ResponseMetadata>
</DescribeReservedCacheNodesOfferingsResponse>
```

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

DescribeServerlessCaches

Returns information about a specific serverless cache. If no identifier is specified, then the API returns information on all the serverless caches belonging to this AWS account.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

MaxResults

The maximum number of records in the response. If more records exist than the specified max-records value, the next token is included in the response so that remaining results can be retrieved. The default is 50.

Type: Integer

Required: No

NextToken

An optional marker returned from a prior request to support pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by MaxResults.

Type: String

Required: No

ServerlessCacheName

The identifier for the serverless cache. If this parameter is specified, only information about that specific serverless cache is returned. Default: NULL

Type: String

Required: No

Response Elements

The following elements are returned by the service.

NextToken

An optional marker returned from a prior request to support pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by MaxResults.

Type: String

ServerlessCaches.member.N

The serverless caches associated with a given description request.

Type: Array of ServerlessCache objects

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

ServerlessCacheNotFoundFault

The serverless cache was not found or does not exist.

HTTP Status Code: 404

Examples

DescribeServerlessCaches

This example illustrates one usage of DescribeServerlessCaches.

Sample Request

```
{
    "input": {
    },
    "output": {
        "ServerlessCaches": [
            {
                "ServerlessCacheName": "my-serverless-cache",
                "Description": "A serverless cache.",
                "Status": "available",
                "Engine": "redis",
                "MajorEngineVersion": "7",
                "FullEngineVersion": "7.0",
                "SubnetIds": [
                    "subnet-xxx8c982",
                    "subnet-xxx382f3",
                    "subnet-xxxb3e7c0"
                ],
                "CacheUsageLimits": {
                    "DataStorage" : {
                        "Maximum" : 10,
                        "Unit" : "GB"
                    },
                    "ECPUPerSecond" : {
                        "Maximum" : 50000
                    }
                },
                "SecurityGroupIds": [
                    "sg-xxx0c9af"
                ],
                "Endpoint": {
                    "Address": "my-serverless-cache-
xxxxxx.serverless.use1qa.cache.amazonaws.com",
                    "Port": 6379
                },
                "ARN": "arn:aws:elasticache:us-east-1:222222222222:serverlesscache:my-
serverless-cache",
                "SnapshotRetentionLimit": 10,
                "DailySnapshotTime": "11:00"
            }
        ]
    }
```

```
}
```

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

DescribeServerlessCacheSnapshots

Returns information about serverless cache snapshots. By default, this API lists all of the customer's serverless cache snapshots. It can also describe a single serverless cache snapshot, or the snapshots associated with a particular serverless cache. Available for Valkey, Redis OSS and Serverless Memcached only.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

MaxResults

The maximum number of records to include in the response. If more records exist than the specified max-results value, a market is included in the response so that remaining results can be retrieved. Available for Valkey, Redis OSS and Serverless Memcached only.The default is 50. The Validation Constraints are a maximum of 50.

Type: Integer

Required: No

NextToken

An optional marker returned from a prior request to support pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by max-results. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: No

ServerlessCacheName

The identifier of serverless cache. If this parameter is specified, only snapshots associated with that specific serverless cache are described. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: No

ServerlessCacheSnapshotName

The identifier of the serverless cache's snapshot. If this parameter is specified, only this snapshot is described. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: No

SnapshotType

The type of snapshot that is being described. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: No

Response Elements

The following elements are returned by the service.

NextToken

An optional marker returned from a prior request to support pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by max-results. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

ServerlessCacheSnapshots.ServerlessCacheSnapshot.N

The serverless caches snapshots associated with a given description request. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: Array of ServerlessCacheSnapshot objects

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

ServerlessCacheNotFoundFault

The serverless cache was not found or does not exist.

HTTP Status Code: 404

ServerlessCacheSnapshotNotFoundFault

This serverless cache snapshot could not be found or does not exist. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 404

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

DescribeServiceUpdates

Returns details of the service updates

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response

Type: Integer

Required: No

ServiceUpdateName

The unique ID of the service update

Type: String

Required: No

ServiceUpdateStatus.member.N

The status of the service update

Type: Array of strings

Array Members: Maximum number of 3 items.

Valid Values: `available | cancelled | expired`

Required: No

Response Elements

The following elements are returned by the service.

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

ServiceUpdates.ServiceUpdate.N

A list of service updates

Type: Array of ServiceUpdate objects

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

ServiceUpdateNotFoundFault

The service update doesn't exist

HTTP Status Code: 404

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

DescribeSnapshots

Returns information about cluster or replication group snapshots. By default, `DescribeSnapshots` lists all of your snapshots; it can optionally describe a single snapshot, or just the snapshots associated with a particular cache cluster.

> **Note**
>
> This operation is valid for Valkey or Redis OSS only.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheClusterId

A user-supplied cluster identifier. If this parameter is specified, only snapshots associated with that specific cluster are described.

Type: String

Required: No

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified `MaxRecords` value, a marker is included in the response so that the remaining results can be retrieved.

Default: 50

Constraints: minimum 20; maximum 50.

Type: Integer

Required: No

ReplicationGroupId

A user-supplied replication group identifier. If this parameter is specified, only snapshots associated with that specific replication group are described.

Type: String

Required: No

ShowNodeGroupConfig

A Boolean value which if true, the node group (shard) configuration is included in the snapshot description.

Type: Boolean

Required: No

SnapshotName

A user-supplied name of the snapshot. If this parameter is specified, only this snapshot are described.

Type: String

Required: No

SnapshotSource

If set to `system`, the output shows snapshots that were automatically created by ElastiCache. If set to `user` the output shows snapshots that were manually created. If omitted, the output shows both automatically and manually created snapshots.

Type: String

Required: No

Response Elements

The following elements are returned by the service.

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Snapshots.Snapshot.N

A list of snapshots. Each item in the list contains detailed information about one snapshot.

Type: Array of Snapshot objects

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterNotFound

The requested cluster ID does not refer to an existing cluster.

HTTP Status Code: 404

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

SnapshotNotFoundFault

The requested snapshot name does not refer to an existing snapshot.

HTTP Status Code: 404

Examples

DescribeSnapshots

This example illustrates one usage of DescribeSnapshots.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=DescribeSnapshots
   &MaxRecords=50
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<DescribeSnapshotsResponse xmlns="http://elasticache.amazonaws.com/doc/2015-02-02/">
   <DescribeSnapshotsResult>
      <Snapshots>
         <Snapshot>
            <CacheClusterId>my-redis-primary</CacheClusterId>
            <Port>6379</Port>
            <CacheNodeType>cache.m1.small</CacheNodeType>
            <CacheParameterGroupName>default.redis2.8</CacheParameterGroupName>
            <Engine>redis</Engine>
            <PreferredAvailabilityZone>us-west-2c</PreferredAvailabilityZone>
            <CacheClusterCreateTime>2015-02-02T18:46:57.972Z</CacheClusterCreateTime>
            <EngineVersion>2.8.6</EngineVersion>
            <SnapshotSource>manual</SnapshotSource>
            <AutoMinorVersionUpgrade>true</AutoMinorVersionUpgrade>
            <PreferredMaintenanceWindow>wed:09:00-wed:10:00</
PreferredMaintenanceWindow>
            <SnapshotName>my-manual-snapshot</SnapshotName>
            <SnapshotRetentionLimit>5</SnapshotRetentionLimit>
            <NodeSnapshots>
               <NodeSnapshot>
                  <SnapshotCreateTime>2015-02-02T18:54:12Z</SnapshotCreateTime>
                  <CacheNodeCreateTime>2015-02-02T18:46:57.972Z</CacheNodeCreateTime>
                  <CacheNodeId>0001</CacheNodeId>
                  <CacheSize>3 MB</CacheSize>
```

```
                </NodeSnapshot>
            </NodeSnapshots>
            <SnapshotStatus>creating</SnapshotStatus>
            <NumCacheNodes>1</NumCacheNodes>
            <SnapshotWindow>07:30-08:30</SnapshotWindow>
         </Snapshot>
      </Snapshots>
   </DescribeSnapshotsResult>
   <ResponseMetadata>
      <RequestId>51b0b25e-b9cf-11e3-8a16-7978bb24ffdf</RequestId>
   </ResponseMetadata>
</DescribeSnapshotsResponse>
```

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

DescribeUpdateActions

Returns details of the update actions

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheClusterIds.member.N

The cache cluster IDs

Type: Array of strings

Array Members: Maximum number of 20 items.

Required: No

Engine

The Elasticache engine to which the update applies. Either Valkey, Redis OSS or Memcached.

Type: String

Required: No

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response

Type: Integer

Required: No

ReplicationGroupIds.member.N

The replication group IDs

Type: Array of strings

Array Members: Maximum number of 20 items.

Required: No

ServiceUpdateName

The unique ID of the service update

Type: String

Required: No

ServiceUpdateStatus.member.N

The status of the service update

Type: Array of strings

Array Members: Maximum number of 3 items.

Valid Values: `available | cancelled | expired`

Required: No

ServiceUpdateTimeRange

The range of time specified to search for service updates that are in available status

Type: TimeRangeFilter object

Required: No

ShowNodeLevelUpdateStatus

Dictates whether to include node level update status in the response

Type: Boolean

Required: No

UpdateActionStatus.member.N

The status of the update action.

Type: Array of strings

Array Members: Maximum number of 9 items.

Valid Values: `not-applied | waiting-to-start | in-progress | stopping | stopped | complete | scheduling | scheduled | not-applicable`

Required: No

Response Elements

The following elements are returned by the service.

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

UpdateActions.UpdateAction.N

Returns a list of update actions

Type: Array of UpdateAction objects

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

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

DescribeUserGroups

Returns a list of user groups.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by MaxRecords. >

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified MaxRecords value, a marker is included in the response so that the remaining results can be retrieved.

Type: Integer

Required: No

UserGroupId

The ID of the user group.

Type: String

Required: No

Response Elements

The following elements are returned by the service.

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by MaxRecords.>

Type: String

UserGroups.member.N

Returns a list of user groups.

Type: Array of UserGroup objects

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

UserGroupNotFound

The user group was not found or does not exist

HTTP Status Code: 404

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

DescribeUsers

Returns a list of users.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

Engine

The engine.

Type: String

Pattern: `[a-zA-Z]*`

Required: No

Filters.member.N

Filter to determine the list of User IDs to return.

Type: Array of Filter objects

Required: No

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by MaxRecords. >

Type: String

Required: No

MaxRecords

The maximum number of records to include in the response. If more records exist than the specified MaxRecords value, a marker is included in the response so that the remaining results can be retrieved.

Type: Integer

Required: No

UserId

The ID of the user.

Type: String

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

Response Elements

The following elements are returned by the service.

Marker

An optional marker returned from a prior request. Use this marker for pagination of results from this operation. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by MaxRecords. >

Type: String

Users.member.N

A list of users.

Type: Array of User objects

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

UserNotFound

The user does not exist or could not be found.

HTTP Status Code: 404

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

DisassociateGlobalReplicationGroup

Remove a secondary cluster from the Global datastore using the Global datastore name. The secondary cluster will no longer receive updates from the primary cluster, but will remain as a standalone cluster in that Amazon region.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

GlobalReplicationGroupId

The name of the Global datastore

Type: String

Required: Yes

ReplicationGroupId

The name of the secondary cluster you wish to remove from the Global datastore

Type: String

Required: Yes

ReplicationGroupRegion

The Amazon region of secondary cluster you wish to remove from the Global datastore

Type: String

Required: Yes

Response Elements

The following element is returned by the service.

GlobalReplicationGroup

Consists of a primary cluster that accepts writes and an associated secondary cluster that resides in a different Amazon region. The secondary cluster accepts only reads. The primary cluster automatically replicates updates to the secondary cluster.

- The **GlobalReplicationGroupIdSuffix** represents the name of the Global datastore, which is what you use to associate a secondary cluster.

Type: GlobalReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

GlobalReplicationGroupNotFoundFault

The Global datastore does not exist

HTTP Status Code: 404

InvalidGlobalReplicationGroupState

The Global datastore is not available or in primary-only state.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

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

ExportServerlessCacheSnapshot

Provides the functionality to export the serverless cache snapshot data to Amazon S3. Available for Valkey and Redis OSS only.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

S3BucketName

Name of the Amazon S3 bucket to export the snapshot to. The Amazon S3 bucket must also be in same region as the snapshot. Available for Valkey and Redis OSS only.

Type: String

Required: Yes

ServerlessCacheSnapshotName

The identifier of the serverless cache snapshot to be exported to S3. Available for Valkey and Redis OSS only.

Type: String

Required: Yes

Response Elements

The following element is returned by the service.

ServerlessCacheSnapshot

The state of a serverless cache at a specific point in time, to the millisecond. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: ServerlessCacheSnapshot object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidServerlessCacheSnapshotStateFault

The state of the serverless cache snapshot was not received. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 400

ServerlessCacheSnapshotNotFoundFault

This serverless cache snapshot could not be found or does not exist. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 404

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

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

FailoverGlobalReplicationGroup

Used to failover the primary region to a secondary region. The secondary region will become primary, and all other clusters will become secondary.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

GlobalReplicationGroupId

The name of the Global datastore

Type: String

Required: Yes

PrimaryRegion

The Amazon region of the primary cluster of the Global datastore

Type: String

Required: Yes

PrimaryReplicationGroupId

The name of the primary replication group

Type: String

Required: Yes

Response Elements

The following element is returned by the service.

GlobalReplicationGroup

Consists of a primary cluster that accepts writes and an associated secondary cluster that resides in a different Amazon region. The secondary cluster accepts only reads. The primary cluster automatically replicates updates to the secondary cluster.

- The **GlobalReplicationGroupIdSuffix** represents the name of the Global datastore, which is what you use to associate a secondary cluster.

Type: GlobalReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

GlobalReplicationGroupNotFoundFault

The Global datastore does not exist

HTTP Status Code: 404

InvalidGlobalReplicationGroupState

The Global datastore is not available or in primary-only state.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

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

IncreaseNodeGroupsInGlobalReplicationGroup

Increase the number of node groups in the Global datastore

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ApplyImmediately

Indicates that the process begins immediately. At present, the only permitted value for this parameter is true.

Type: Boolean

Required: Yes

GlobalReplicationGroupId

The name of the Global datastore

Type: String

Required: Yes

NodeGroupCount

Total number of node groups you want

Type: Integer

Required: Yes

RegionalConfigurations.RegionalConfiguration.N

Describes the replication group IDs, the Amazon regions where they are stored and the shard configuration for each that comprise the Global datastore

Type: Array of RegionalConfiguration objects

Required: No

Response Elements

The following element is returned by the service.

GlobalReplicationGroup

Consists of a primary cluster that accepts writes and an associated secondary cluster that resides in a different Amazon region. The secondary cluster accepts only reads. The primary cluster automatically replicates updates to the secondary cluster.

- The **GlobalReplicationGroupIdSuffix** represents the name of the Global datastore, which is what you use to associate a secondary cluster.

Type: GlobalReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

GlobalReplicationGroupNotFoundFault

The Global datastore does not exist

HTTP Status Code: 404

InvalidGlobalReplicationGroupState

The Global datastore is not available or in primary-only state.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

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

IncreaseReplicaCount

Dynamically increases the number of replicas in a Valkey or Redis OSS (cluster mode disabled) replication group or the number of replica nodes in one or more node groups (shards) of a Valkey or Redis OSS (cluster mode enabled) replication group. This operation is performed with no cluster down time.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ApplyImmediately

If `True`, the number of replica nodes is increased immediately. `ApplyImmediately=False` is not currently supported.

Type: Boolean

Required: Yes

ReplicationGroupId

The id of the replication group to which you want to add replica nodes.

Type: String

Required: Yes

NewReplicaCount

The number of read replica nodes you want at the completion of this operation. For Valkey or Redis OSS (cluster mode disabled) replication groups, this is the number of replica nodes in the replication group. For Valkey or Redis OSS (cluster mode enabled) replication groups, this is the number of replica nodes in each of the replication group's node groups.

Type: Integer

Required: No

ReplicaConfiguration.ConfigureShard.N

A list of `ConfigureShard` objects that can be used to configure each shard in a Valkey or Redis OSS (cluster mode enabled) replication group. The `ConfigureShard` has three members: `NewReplicaCount`, `NodeGroupId`, and `PreferredAvailabilityZones`.

Type: Array of ConfigureShard objects

Required: No

Response Elements

The following element is returned by the service.

ReplicationGroup

Contains all of the attributes of a specific Valkey or Redis OSS replication group.

Type: ReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

ClusterQuotaForCustomerExceeded

The request cannot be processed because it would exceed the allowed number of clusters per customer.

HTTP Status Code: 400

InsufficientCacheClusterCapacity

The requested cache node type is not available in the specified Availability Zone. For more information, see InsufficientCacheClusterCapacity in the ElastiCache User Guide.

HTTP Status Code: 400

InvalidCacheClusterState

The requested cluster is not in the `available` state.

HTTP Status Code: 400

InvalidKMSKeyFault

The KMS key supplied is not valid.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

InvalidVPCNetworkStateFault

The VPC network is in an invalid state.

HTTP Status Code: 400

NodeGroupsPerReplicationGroupQuotaExceeded

The request cannot be processed because it would exceed the maximum allowed number of node groups (shards) in a single replication group. The default maximum is 90

HTTP Status Code: 400

NodeQuotaForCustomerExceeded

The request cannot be processed because it would exceed the allowed number of cache nodes per customer.

HTTP Status Code: 400

NoOperationFault

The operation was not performed because no changes were required.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

Examples

Example

The following example increases the replica count to 3 in all node groups of `sample-repl-group`.

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=IncreaseReplicaCount
   &ApplyImmediately=True
   &NewReplicaCount=3
   &ReplicationGroupId=sample-repl-group
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Example

The following example increases the replica count in two node groups. Because there are multiple node groups, this code only works for Valkey or Redis OSS (cluster mode enabled) replication groups.

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=IncreaseReplicaCount
   &ApplyImmediately=True
   &ReplicaConfiguration.ConfigureShard.1.NodeGroupId=0001
   &ReplicaConfiguration.ConfigureShard.1.NewReplicaCount=2

 &ReplicaConfiguration.ConfigureShard.1.PreferredAvailabilityZones.PreferredAvailabilityZone.1=
east-1a

 &ReplicaConfiguration.ConfigureShard.1.PreferredAvailabilityZones.PreferredAvailabilityZone.2=
east-1c

 &ReplicaConfiguration.ConfigureShard.1.PreferredAvailabilityZones.PreferredAvailabilityZone.3=
east-1b
   &ReplicaConfiguration.ConfigureShard.2.NodeGroupId=0003
   &ReplicaConfiguration.ConfigureShard.2.NewReplicaCount=3
```

```
 &ReplicaConfiguration.ConfigureShard.2.PreferredAvailabilityZones.PreferredAvailabilityZone.1=
east-1a

 &ReplicaConfiguration.ConfigureShard.2.PreferredAvailabilityZones.PreferredAvailabilityZone.2=
east-1b

 &ReplicaConfiguration.ConfigureShard.2.PreferredAvailabilityZones.PreferredAvailabilityZone.3=
east-1c

 &ReplicaConfiguration.ConfigureShard.2.PreferredAvailabilityZones.PreferredAvailabilityZone.4=
east-1c
   &ReplicationGroupId=sample-repl-group
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

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

ListAllowedNodeTypeModifications

Lists all available node types that you can scale with your cluster's replication group's current node type.

When you use the `ModifyCacheCluster` or `ModifyReplicationGroup` operations to scale your cluster or replication group, the value of the `CacheNodeType` parameter must be one of the node types returned by this operation.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheClusterId

The name of the cluster you want to scale up to a larger node instanced type. ElastiCache uses the cluster id to identify the current node type of this cluster and from that to create a list of node types you can scale up to.

> **Important**
>
> You must provide a value for either the `CacheClusterId` or the `ReplicationGroupId`.

Type: String

Required: No

ReplicationGroupId

The name of the replication group want to scale up to a larger node type. ElastiCache uses the replication group id to identify the current node type being used by this replication group, and from that to create a list of node types you can scale up to.

> **Important**
>
> You must provide a value for either the `CacheClusterId` or the `ReplicationGroupId`.

Type: String

Required: No

Response Elements

The following elements are returned by the service.

ScaleDownModifications.member.N

A string list, each element of which specifies a cache node type which you can use to scale your cluster or replication group. When scaling down a Valkey or Redis OSS cluster or replication group using ModifyCacheCluster or ModifyReplicationGroup, use a value from this list for the CacheNodeType parameter.

Type: Array of strings

ScaleUpModifications.member.N

A string list, each element of which specifies a cache node type which you can use to scale your cluster or replication group.

When scaling up a Valkey or Redis OSS cluster or replication group using `ModifyCacheCluster` or `ModifyReplicationGroup`, use a value from this list for the `CacheNodeType` parameter.

Type: Array of strings

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterNotFound

The requested cluster ID does not refer to an existing cluster.

HTTP Status Code: 404

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

Examples

ListAllowedNodeTypeModifications for a Cluster

The following example request a list of node types you can use to scale `myCluster` up.

Sample Request

```
https://elasticache.us-east-1.amazonaws.com/
   ?Action=ListAllowedNodeTypeModifications
   &CacheClusterId=mycachecluster
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Version=2015-02-02
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

ListAllowedNodeTypeModifications for a Replication Group

The following example requests a list of node types you can use to scale `myReplGroup` up.

Sample Request

```
https://elasticache.us-east-1.amazonaws.com/
   ?Action=ListAllowedNodeTypeModifications
   &ReplicationGroupId=myreplgroup
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Version=2015-02-02
```

```
    &Timestamp=20150202T192317Z
    &X-Amz-Credential=<credential>
```

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

Lists all tags currently on a named resource.

A tag is a key-value pair where the key and value are case-sensitive. You can use tags to categorize and track all your ElastiCache resources, with the exception of global replication group. When you add or remove tags on replication groups, those actions will be replicated to all nodes in the replication group. For more information, see Resource-level permissions.

If the cluster is not in the *available* state, `ListTagsForResource` returns an error.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ResourceName

The Amazon Resource Name (ARN) of the resource for which you want the list of tags, for example `arn:aws:elasticache:us-west-2:0123456789:cluster:myCluster` or `arn:aws:elasticache:us-west-2:0123456789:snapshot:mySnapshot`.

For more information about ARNs, see Amazon Resource Names (ARNs) and AWS Service Namespaces.

Type: String

Required: Yes

Response Elements

The following element is returned by the service.

TagList.Tag.N

A list of tags as key-value pairs.

Type: Array of Tag objects

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterNotFound

The requested cluster ID does not refer to an existing cluster.

HTTP Status Code: 404

CacheParameterGroupNotFound

The requested cache parameter group name does not refer to an existing cache parameter group.

HTTP Status Code: 404

CacheSecurityGroupNotFound

The requested cache security group name does not refer to an existing cache security group.

HTTP Status Code: 404

CacheSubnetGroupNotFoundFault

The requested cache subnet group name does not refer to an existing cache subnet group.

HTTP Status Code: 400

InvalidARN

The requested Amazon Resource Name (ARN) does not refer to an existing resource.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

InvalidServerlessCacheSnapshotStateFault

The state of the serverless cache snapshot was not received. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 400

InvalidServerlessCacheStateFault

The account for these credentials is not currently active.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

ReservedCacheNodeNotFound

The requested reserved cache node was not found.

HTTP Status Code: 404

ServerlessCacheNotFoundFault

The serverless cache was not found or does not exist.

HTTP Status Code: 404

ServerlessCacheSnapshotNotFoundFault

This serverless cache snapshot could not be found or does not exist. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 404

SnapshotNotFoundFault

The requested snapshot name does not refer to an existing snapshot.

HTTP Status Code: 404

UserGroupNotFound

The user group was not found or does not exist

HTTP Status Code: 404

UserNotFound

The user does not exist or could not be found.

HTTP Status Code: 404

Examples

ListTagsForResource

This example illustrates one usage of ListTagsForResource.

Sample Request

```
https://elasticache.us-east-1.amazonaws.com/
   ?Action=ListTagsForResource
   &ResourceName=arn:aws:elasticache:us-west-2:0123456789:cluster:myCluster
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Version=2015-02-02
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

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

ModifyCacheCluster

Modifies the settings for a cluster. You can use this operation to change one or more cluster configuration parameters by specifying the parameters and the new values.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheClusterId

The cluster identifier. This value is stored as a lowercase string.

Type: String

Required: Yes

ApplyImmediately

If `true`, this parameter causes the modifications in this request and any pending modifications to be applied, asynchronously and as soon as possible, regardless of the `PreferredMaintenanceWindow` setting for the cluster.

If `false`, changes to the cluster are applied on the next maintenance reboot, or the next failure reboot, whichever occurs first.

> **Important**
>
> If you perform a `ModifyCacheCluster` before a pending modification is applied, the pending modification is replaced by the newer modification.

Valid values: `true` | `false`

Default: `false`

Type: Boolean

Required: No

AuthToken

Reserved parameter. The password used to access a password protected server. This parameter must be specified with the `auth-token-update` parameter. Password constraints:

- Must be only printable ASCII characters
- Must be at least 16 characters and no more than 128 characters in length
- Cannot contain any of the following characters: '/', '"', or '@', '%'

For more information, see AUTH password at AUTH.

Type: String

Required: No

AuthTokenUpdateStrategy

Specifies the strategy to use to update the AUTH token. This parameter must be specified with the `auth-token` parameter. Possible values:

- ROTATE - default, if no update strategy is provided
- SET - allowed only after ROTATE
- DELETE - allowed only when transitioning to RBAC

For more information, see Authenticating Users with AUTH

Type: String

Valid Values: `SET | ROTATE | DELETE`

Required: No

AutoMinorVersionUpgrade

If you are running Valkey 7.2 or Redis OSS engine version 6.0 or later, set this parameter to yes to opt-in to the next auto minor version upgrade campaign. This parameter is disabled for previous versions.

Type: Boolean

Required: No

AZMode

Specifies whether the new nodes in this Memcached cluster are all created in a single Availability Zone or created across multiple Availability Zones.

Valid values: `single-az` | `cross-az`.

This option is only supported for Memcached clusters.

> **Note**
>
> You cannot specify `single-az` if the Memcached cluster already has cache nodes in different Availability Zones. If `cross-az` is specified, existing Memcached nodes remain in their current Availability Zone.
> Only newly created nodes are located in different Availability Zones.

Type: String

Valid Values: `single-az | cross-az`

Required: No

CacheNodeIdsToRemove.CacheNodeId.N

A list of cache node IDs to be removed. A node ID is a numeric identifier (0001, 0002, etc.). This parameter is only valid when `NumCacheNodes` is less than the existing number of cache nodes. The number of cache node IDs supplied in this parameter must match the difference between the existing number of cache nodes in the cluster or pending cache nodes, whichever is greater, and the value of `NumCacheNodes` in the request.

For example: If you have 3 active cache nodes, 7 pending cache nodes, and the number of cache nodes in this `ModifyCacheCluster` call is 5, you must list 2 (7 - 5) cache node IDs to remove.

Type: Array of strings

Required: No

CacheNodeType

A valid cache node type that you want to scale this cluster up to.

Type: String

Required: No

CacheParameterGroupName

The name of the cache parameter group to apply to this cluster. This change is asynchronously applied as soon as possible for parameters when the `ApplyImmediately` parameter is specified as `true` for this request.

Type: String

Required: No

CacheSecurityGroupNames.CacheSecurityGroupName.N

A list of cache security group names to authorize on this cluster. This change is asynchronously applied as soon as possible.

You can use this parameter only with clusters that are created outside of an Amazon Virtual Private Cloud (Amazon VPC).

Constraints: Must contain no more than 255 alphanumeric characters. Must not be "Default".

Type: Array of strings

Required: No

Engine

The engine type used by the cache cluster. The options are valkey, memcached or redis.

Type: String

Required: No

EngineVersion

The upgraded version of the cache engine to be run on the cache nodes.

Important: You can upgrade to a newer engine version (see Selecting a Cache Engine and Version), but you cannot downgrade to an earlier engine version. If you want to use an earlier engine version, you must delete the existing cluster and create it anew with the earlier engine version.

Type: String

Required: No

IpDiscovery

The network type you choose when modifying a cluster, either `ipv4` | `ipv6`. IPv6 is supported for workloads using Valkey 7.2 and above, Redis OSS engine version 6.2 to 7.1 or Memcached engine version 1.6.6 and above on all instances built on the Nitro system.

Type: String

Valid Values: `ipv4 | ipv6`

Required: No

LogDeliveryConfigurations.LogDeliveryConfigurationRequest.N

Specifies the destination, format and type of the logs.

Type: Array of LogDeliveryConfigurationRequest objects

Required: No

NewAvailabilityZones.PreferredAvailabilityZone.N

> **Note**
>
> This option is only supported on Memcached clusters.

The list of Availability Zones where the new Memcached cache nodes are created.

This parameter is only valid when `NumCacheNodes` in the request is greater than the sum of the number of active cache nodes and the number of cache nodes pending creation (which may be zero). The number of Availability Zones supplied in this list must match the cache nodes being added in this request.

Scenarios:

- **Scenario 1:** You have 3 active nodes and wish to add 2 nodes. Specify `NumCacheNodes=5` (3 + 2) and optionally specify two Availability Zones for the two new nodes.
- **Scenario 2:** You have 3 active nodes and 2 nodes pending creation (from the scenario 1 call) and want to add 1 more node. Specify `NumCacheNodes=6` ((3 + 2) + 1) and optionally specify an Availability Zone for the new node.
- **Scenario 3:** You want to cancel all pending operations. Specify `NumCacheNodes=3` to cancel all pending operations.

The Availability Zone placement of nodes pending creation cannot be modified. If you wish to cancel any nodes pending creation, add 0 nodes by setting `NumCacheNodes` to the number of current nodes.

If `cross-az` is specified, existing Memcached nodes remain in their current Availability Zone. Only newly created nodes can be located in different Availability Zones. For guidance on how to move existing Memcached nodes to different Availability Zones, see the **Availability Zone Considerations** section of Cache Node Considerations for Memcached.

Impact of new add/remove requests upon pending requests

- Scenario-1
  - Pending Action: Delete
  - New Request: Delete
  - Result: The new delete, pending or immediate, replaces the pending delete.
- Scenario-2
  - Pending Action: Delete
  - New Request: Create
  - Result: The new create, pending or immediate, replaces the pending delete.
- Scenario-3
  - Pending Action: Create
  - New Request: Delete
  - Result: The new delete, pending or immediate, replaces the pending create.
- Scenario-4
  - Pending Action: Create
  - New Request: Create
  - Result: The new create is added to the pending create.

> **Important**
>
> **Important:** If the new create request is **Apply Immediately - Yes**, all creates are performed immediately. If the new create request is **Apply Immediately - No**, all creates are pending.

Type: Array of strings

Required: No

NotificationTopicArn

The Amazon Resource Name (ARN) of the Amazon SNS topic to which notifications are sent.

> **Note**
>
> The Amazon SNS topic owner must be same as the cluster owner.

Type: String

Required: No

NotificationTopicStatus

The status of the Amazon SNS notification topic. Notifications are sent only if the status is `active`.

Valid values: `active` | `inactive`

Type: String

Required: No

NumCacheNodes

The number of cache nodes that the cluster should have. If the value for `NumCacheNodes` is greater than the sum of the number of current cache nodes and the number of cache nodes pending creation (which may be zero), more nodes are added. If the value is less than the number of existing cache nodes, nodes are removed. If the value is equal to the number of current cache nodes, any pending add or remove requests are canceled.

If you are removing cache nodes, you must use the `CacheNodeIdsToRemove` parameter to provide the IDs of the specific cache nodes to remove.

For clusters running Valkey or Redis OSS, this value must be 1. For clusters running Memcached, this value must be between 1 and 40.

> **Note**
>
> Adding or removing Memcached cache nodes can be applied immediately or as a pending operation (see `ApplyImmediately`).
> A pending operation to modify the number of cache nodes in a cluster during its maintenance window, whether by adding or removing nodes in accordance with the scale out architecture, is not queued. The customer's latest request to add or remove nodes to the cluster overrides any previous pending operations to modify the number of cache nodes in the cluster. For example, a request to remove 2 nodes would override a previous pending operation to remove 3 nodes. Similarly, a request to add 2 nodes would override a previous pending operation to remove 3 nodes and vice versa. As Memcached cache nodes may now be provisioned in different Availability Zones with flexible cache node placement, a request to add nodes does not automatically override

a previous pending operation to add nodes. The customer can modify the previous pending operation to add more nodes or explicitly cancel the pending request and retry the new request. To cancel pending operations to modify the number of cache nodes in a cluster, use the `ModifyCacheCluster` request and set `NumCacheNodes` equal to the number of cache nodes currently in the cluster.

Type: Integer

Required: No

PreferredMaintenanceWindow

Specifies the weekly time range during which maintenance on the cluster is performed. It is specified as a range in the format ddd:hh24:mi-ddd:hh24:mi (24H Clock UTC). The minimum maintenance window is a 60 minute period.

Valid values for ddd are:

- `sun`
- `mon`
- `tue`
- `wed`
- `thu`
- `fri`
- `sat`

Example: `sun:23:00-mon:01:30`

Type: String

Required: No

ScaleConfig

Configures horizontal or vertical scaling for Memcached clusters, specifying the scaling percentage and interval.

Type: ScaleConfig object

Required: No

SecurityGroupIds.SecurityGroupId.N

Specifies the VPC Security Groups associated with the cluster.

This parameter can be used only with clusters that are created in an Amazon Virtual Private Cloud (Amazon VPC).

Type: Array of strings

Required: No

SnapshotRetentionLimit

The number of days for which ElastiCache retains automatic cluster snapshots before deleting them. For example, if you set `SnapshotRetentionLimit` to 5, a snapshot that was taken today is retained for 5 days before being deleted.

> **Note**
>
> If the value of `SnapshotRetentionLimit` is set to zero (0), backups are turned off.

Type: Integer

Required: No

SnapshotWindow

The daily time range (in UTC) during which ElastiCache begins taking a daily snapshot of your cluster.

Type: String

Required: No

Response Elements

The following element is returned by the service.

CacheCluster

Contains all of the attributes of a specific cluster.

Type: CacheCluster object

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterNotFound

The requested cluster ID does not refer to an existing cluster.

HTTP Status Code: 404

CacheParameterGroupNotFound

The requested cache parameter group name does not refer to an existing cache parameter group.

HTTP Status Code: 404

CacheSecurityGroupNotFound

The requested cache security group name does not refer to an existing cache security group.

HTTP Status Code: 404

InsufficientCacheClusterCapacity

The requested cache node type is not available in the specified Availability Zone. For more information, see InsufficientCacheClusterCapacity in the ElastiCache User Guide.

HTTP Status Code: 400

InvalidCacheClusterState

The requested cluster is not in the `available` state.

HTTP Status Code: 400

InvalidCacheSecurityGroupState

The current state of the cache security group does not allow deletion.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidVPCNetworkStateFault

The VPC network is in an invalid state.

HTTP Status Code: 400

NodeQuotaForClusterExceeded

The request cannot be processed because it would exceed the allowed number of cache nodes in a single cluster.

HTTP Status Code: 400

NodeQuotaForCustomerExceeded

The request cannot be processed because it would exceed the allowed number of cache nodes per customer.

HTTP Status Code: 400

Examples

ModifyCacheCluster

This example illustrates one usage of ModifyCacheCluster.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=ModifyCacheCluster
   NumCacheNodes=5
   &CacheClusterId=simcoprod01
   &ApplyImmediately=true
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
```

```
    &Timestamp=20150202T192317Z
    &X-Amz-Credential=<credential>
```

Sample Response

```
<ModifyCacheClusterResponse xmlns="http://elasticache.amazonaws.com/doc/2015-02-02/">
  <ModifyCacheClusterResult>
    <CacheCluster>
      <CacheParameterGroup>
        <ParameterApplyStatus>in-sync</ParameterApplyStatus>
        <CacheParameterGroupName>default.memcached1.4</CacheParameterGroupName>
        <CacheNodeIdsToReboot/>
      </CacheParameterGroup>
      <CacheClusterId>simcoprod01</CacheClusterId>
      <CacheClusterStatus>available</CacheClusterStatus>
      <ConfigurationEndpoint>
        <Port>11211</Port>
        <Address>simcoprod01.m2st2p.cfg.cache.amazonaws.com</Address>
      </ConfigurationEndpoint>
      <CacheNodeType>cache.m1.large</CacheNodeType>
      <Engine>memcached</Engine>
      <PendingModifiedValues>
        <NumCacheNodes>5</NumCacheNodes>
      </PendingModifiedValues>
      <PreferredAvailabilityZone>us-west-2b</PreferredAvailabilityZone>
      <CacheClusterCreateTime>2015-02-02T23:45:20.937Z</CacheClusterCreateTime>
      <EngineVersion>1.4.5</EngineVersion>
      <AutoMinorVersionUpgrade>true</AutoMinorVersionUpgrade>
      <PreferredMaintenanceWindow>fri:04:30-fri:05:30</PreferredMaintenanceWindow>
      <CacheSecurityGroups>
        <CacheSecurityGroup>
          <CacheSecurityGroupName>default</CacheSecurityGroupName>
          <Status>active</Status>
        </CacheSecurityGroup>
      </CacheSecurityGroups>
      <NumCacheNodes>3</NumCacheNodes>
    </CacheCluster>
  </ModifyCacheClusterResult>
  <ResponseMetadata>
    <RequestId>d5786c6d-b7fe-11e0-9326-b7275b9d4a6c</RequestId>
  </ResponseMetadata>
</ModifyCacheClusterResponse>
```

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

ModifyCacheParameterGroup

Modifies the parameters of a cache parameter group. You can modify up to 20 parameters in a single request by submitting a list parameter name and value pairs.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheParameterGroupName

The name of the cache parameter group to modify.

Type: String

Required: Yes

ParameterNameValues.ParameterNameValue.N

An array of parameter names and values for the parameter update. You must supply at least one parameter name and value; subsequent arguments are optional. A maximum of 20 parameters may be modified per request.

Type: Array of ParameterNameValue objects

Required: Yes

Response Elements

The following element is returned by the service.

CacheParameterGroupName

The name of the cache parameter group.

Type: String

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheParameterGroupNotFound

The requested cache parameter group name does not refer to an existing cache parameter group.

HTTP Status Code: 404

InvalidCacheParameterGroupState

The current state of the cache parameter group does not allow the requested operation to occur.

HTTP Status Code: 400

InvalidGlobalReplicationGroupState

The Global datastore is not available or in primary-only state.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

Examples

ModifyCacheParameterGroup

This example illustrates one usage of ModifyCacheParameterGroup.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=ModifyCacheParameterGroup
   &CacheParameterGroupName=mycacheparametergroup
   &ParameterNameValues.ParameterNameValue.1.ParameterName=chunk_size_growth_factor
```

```
    &ParameterNameValues.ParameterNameValue.1.ParameterValue=1.02
    &Version=2015-02-02
    &SignatureVersion=4
    &SignatureMethod=HmacSHA256
    &Timestamp=20150202T192317Z
    &X-Amz-Credential=<credential>
```

Sample Response

```
<ModifyCacheParameterGroupResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
  <ModifyCacheParameterGroupResult>
    <CacheParameterGroupName>mycacheparametergroup</CacheParameterGroupName>
  </ModifyCacheParameterGroupResult>
  <ResponseMetadata>
    <RequestId>fcedeef2-b7ff-11e0-9326-b7275b9d4a6c</RequestId>
  </ResponseMetadata>
</ModifyCacheParameterGroupResponse>
```

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

ModifyCacheSubnetGroup

Modifies an existing cache subnet group.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheSubnetGroupName

The name for the cache subnet group. This value is stored as a lowercase string.

Constraints: Must contain no more than 255 alphanumeric characters or hyphens.

Example: `mysubnetgroup`

Type: String

Required: Yes

CacheSubnetGroupDescription

A description of the cache subnet group.

Type: String

Required: No

SubnetIds.SubnetIdentifier.N

The EC2 subnet IDs for the cache subnet group.

Type: Array of strings

Required: No

Response Elements

The following element is returned by the service.

CacheSubnetGroup

Represents the output of one of the following operations:

- `CreateCacheSubnetGroup`
- `ModifyCacheSubnetGroup`

Type: CacheSubnetGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheSubnetGroupNotFoundFault

The requested cache subnet group name does not refer to an existing cache subnet group.

HTTP Status Code: 400

CacheSubnetQuotaExceededFault

The request cannot be processed because it would exceed the allowed number of subnets in a cache subnet group.

HTTP Status Code: 400

InvalidSubnet

An invalid subnet identifier was specified.

HTTP Status Code: 400

SubnetInUse

The requested subnet is being used by another cache subnet group.

HTTP Status Code: 400

SubnetNotAllowedFault

At least one subnet ID does not match the other subnet IDs. This mismatch typically occurs when a user sets one subnet ID to a regional Availability Zone and a different one to an outpost. Or when a user sets the subnet ID to an Outpost when not subscribed on this service.

HTTP Status Code: 400

Examples

ModifyCacheSubnetGroup

This example illustrates one usage of ModifyCacheSubnetGroup.

Sample Request

```
https://elasticache.amazonaws.com/
   ?Action=ModifyCacheSubnetGroup
   &CacheSubnetGroupName=myCachesubnetgroup
   &CacheSubnetGroupDescription=My%20modified%20CacheSubnetGroup
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<ModifyCacheSubnetGroupResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
    <ModifyCacheSubnetGroupResult>
            <CacheSubnetGroup>
                <VpcId>990524496922</VpcId>
                <CacheSubnetGroupDescription>My modified CacheSubnetGroup</
CacheSubnetGroupDescription>
                <CacheSubnetGroupName>myCachesubnetgroup</CacheSubnetGroupName>
                <Subnets>
                    <Subnet>
                        <SubnetStatus>Active</SubnetStatus>
                        <SubnetIdentifier>subnet-7c5b4115</SubnetIdentifier>
                        <SubnetAvailabilityZone>
                            <Name>us-west-2c</Name>
                        </SubnetAvailabilityZone>
                    </Subnet>
                    <Subnet>
                        <SubnetStatus>Active</SubnetStatus>
                        <SubnetIdentifier>subnet-7b5b4112</SubnetIdentifier>
                        <SubnetAvailabilityZone>
                            <Name>us-west-2b</Name>
                        </SubnetAvailabilityZone>
                    </Subnet>
```

```
                    <Subnet>
                        <SubnetStatus>Active</SubnetStatus>
                        <SubnetIdentifier>subnet-3ea6bd57</SubnetIdentifier>
                        <SubnetAvailabilityZone>
                            <Name>us-west-2c</Name>
                        </SubnetAvailabilityZone>
                    </Subnet>
                </Subnets>
            </CacheSubnetGroup>
        </ModifyCacheSubnetGroupResult>
        <ResponseMetadata>
            <RequestId>ed662948-a57b-11df-9e38-7ffab86c801f</RequestId>
        </ResponseMetadata>
    </ModifyCacheSubnetGroupResponse>
```

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

ModifyGlobalReplicationGroup

Modifies the settings for a Global datastore.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ApplyImmediately

This parameter causes the modifications in this request and any pending modifications to be applied, asynchronously and as soon as possible. Modifications to Global Replication Groups cannot be requested to be applied in PreferredMaintenceWindow.

Type: Boolean

Required: Yes

GlobalReplicationGroupId

The name of the Global datastore

Type: String

Required: Yes

AutomaticFailoverEnabled

Determines whether a read replica is automatically promoted to read/write primary if the existing primary encounters a failure.

Type: Boolean

Required: No

CacheNodeType

A valid cache node type that you want to scale this Global datastore to.

Type: String

Required: No

CacheParameterGroupName

The name of the cache parameter group to use with the Global datastore. It must be compatible with the major engine version used by the Global datastore.

Type: String

Required: No

Engine

Modifies the engine listed in a global replication group message. The options are valkey, memcached or redis.

Type: String

Required: No

EngineVersion

The upgraded version of the cache engine to be run on the clusters in the Global datastore.

Type: String

Required: No

GlobalReplicationGroupDescription

A description of the Global datastore

Type: String

Required: No

Response Elements

The following element is returned by the service.

GlobalReplicationGroup

Consists of a primary cluster that accepts writes and an associated secondary cluster that resides in a different Amazon region. The secondary cluster accepts only reads. The primary cluster automatically replicates updates to the secondary cluster.

- The **GlobalReplicationGroupIdSuffix** represents the name of the Global datastore, which is what you use to associate a secondary cluster.

Type: GlobalReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

GlobalReplicationGroupNotFoundFault

The Global datastore does not exist

HTTP Status Code: 404

InvalidGlobalReplicationGroupState

The Global datastore is not available or in primary-only state.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

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

ModifyReplicationGroup

Modifies the settings for a replication group. This is limited to Valkey and Redis OSS 7 and above.

- Scaling for Valkey or Redis OSS (cluster mode enabled) in the ElastiCache User Guide
- ModifyReplicationGroupShardConfiguration in the ElastiCache API Reference

> **Note**
>
> This operation is valid for Valkey or Redis OSS only.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ReplicationGroupId

The identifier of the replication group to modify.

Type: String

Required: Yes

ApplyImmediately

If `true`, this parameter causes the modifications in this request and any pending modifications to be applied, asynchronously and as soon as possible, regardless of the `PreferredMaintenanceWindow` setting for the replication group.

If `false`, changes to the nodes in the replication group are applied on the next maintenance reboot, or the next failure reboot, whichever occurs first.

Valid values: `true` | `false`

Default: `false`

Type: Boolean

Required: No

AuthToken

Reserved parameter. The password used to access a password protected server. This parameter must be specified with the `auth-token-update-strategy` parameter. Password constraints:

- Must be only printable ASCII characters
- Must be at least 16 characters and no more than 128 characters in length
- Cannot contain any of the following characters: '/', '"', or '@', '%'

For more information, see AUTH password at AUTH.

Type: String

Required: No

AuthTokenUpdateStrategy

Specifies the strategy to use to update the AUTH token. This parameter must be specified with the `auth-token` parameter. Possible values:

- ROTATE - default, if no update strategy is provided
- SET - allowed only after ROTATE
- DELETE - allowed only when transitioning to RBAC

For more information, see Authenticating Users with AUTH

Type: String

Valid Values: `SET` | `ROTATE` | `DELETE`

Required: No

AutomaticFailoverEnabled

Determines whether a read replica is automatically promoted to read/write primary if the existing primary encounters a failure.

Valid values: `true` | `false`

Type: Boolean

Required: No

AutoMinorVersionUpgrade

If you are running Valkey or Redis OSS engine version 6.0 or later, set this parameter to yes if you want to opt-in to the next auto minor version upgrade campaign. This parameter is disabled for previous versions.

Type: Boolean

Required: No

CacheNodeType

A valid cache node type that you want to scale this replication group to.

Type: String

Required: No

CacheParameterGroupName

The name of the cache parameter group to apply to all of the clusters in this replication group. This change is asynchronously applied as soon as possible for parameters when the `ApplyImmediately` parameter is specified as `true` for this request.

Type: String

Required: No

CacheSecurityGroupNames.CacheSecurityGroupName.N

A list of cache security group names to authorize for the clusters in this replication group. This change is asynchronously applied as soon as possible.

This parameter can be used only with replication group containing clusters running outside of an Amazon Virtual Private Cloud (Amazon VPC).

Constraints: Must contain no more than 255 alphanumeric characters. Must not be `Default`.

Type: Array of strings

Required: No

ClusterMode

Enabled or Disabled. To modify cluster mode from Disabled to Enabled, you must first set the cluster mode to Compatible. Compatible mode allows your Valkey or Redis OSS clients

to connect using both cluster mode enabled and cluster mode disabled. After you migrate all Valkey or Redis OSS clients to use cluster mode enabled, you can then complete cluster mode configuration and set the cluster mode to Enabled.

Type: String

Valid Values: `enabled | disabled | compatible`

Required: No

Engine

Modifies the engine listed in a replication group message. The options are valkey, memcached or redis.

Type: String

Required: No

EngineVersion

The upgraded version of the cache engine to be run on the clusters in the replication group.

Important: You can upgrade to a newer engine version (see Selecting a Cache Engine and Version), but you cannot downgrade to an earlier engine version. If you want to use an earlier engine version, you must delete the existing replication group and create it anew with the earlier engine version.

Type: String

Required: No

IpDiscovery

The network type you choose when modifying a cluster, either `ipv4` | `ipv6`. IPv6 is supported for workloads using Valkey 7.2 and above, Redis OSS engine version 6.2 to 7.1 and Memcached engine version 1.6.6 and above on all instances built on the Nitro system.

Type: String

Valid Values: `ipv4 | ipv6`

Required: No

LogDeliveryConfigurations.LogDeliveryConfigurationRequest.N

Specifies the destination, format and type of the logs.

Type: Array of LogDeliveryConfigurationRequest objects

Required: No

MultiAZEnabled

A flag to indicate MultiAZ is enabled.

Type: Boolean

Required: No

NodeGroupId

Deprecated. This parameter is not used.

Type: String

Required: No

NotificationTopicArn

The Amazon Resource Name (ARN) of the Amazon SNS topic to which notifications are sent.

> **Note**
>
> The Amazon SNS topic owner must be same as the replication group owner.

Type: String

Required: No

NotificationTopicStatus

The status of the Amazon SNS notification topic for the replication group. Notifications are sent only if the status is `active`.

Valid values: `active` | `inactive`

Type: String

Required: No

PreferredMaintenanceWindow

Specifies the weekly time range during which maintenance on the cluster is performed. It is specified as a range in the format ddd:hh24:mi-ddd:hh24:mi (24H Clock UTC). The minimum maintenance window is a 60 minute period.

Valid values for ddd are:

- `sun`
- `mon`
- `tue`
- `wed`
- `thu`
- `fri`
- `sat`

Example: `sun:23:00-mon:01:30`

Type: String

Required: No

PrimaryClusterId

For replication groups with a single primary, if this parameter is specified, ElastiCache promotes the specified cluster in the specified replication group to the primary role. The nodes of all other clusters in the replication group are read replicas.

Type: String

Required: No

RemoveUserGroups

Removes the user group associated with this replication group.

Type: Boolean

Required: No

ReplicationGroupDescription

A description for the replication group. Maximum length is 255 characters.

Type: String

Required: No

SecurityGroupIds.SecurityGroupId.N

Specifies the VPC Security Groups associated with the clusters in the replication group.

This parameter can be used only with replication group containing clusters running in an Amazon Virtual Private Cloud (Amazon VPC).

Type: Array of strings

Required: No

SnapshotRetentionLimit

The number of days for which ElastiCache retains automatic node group (shard) snapshots before deleting them. For example, if you set `SnapshotRetentionLimit` to 5, a snapshot that was taken today is retained for 5 days before being deleted.

Important If the value of SnapshotRetentionLimit is set to zero (0), backups are turned off.

Type: Integer

Required: No

SnapshottingClusterId

The cluster ID that is used as the daily snapshot source for the replication group. This parameter cannot be set for Valkey or Redis OSS (cluster mode enabled) replication groups.

Type: String

Required: No

SnapshotWindow

The daily time range (in UTC) during which ElastiCache begins taking a daily snapshot of the node group (shard) specified by `SnapshottingClusterId`.

Example: `05:00-09:00`

If you do not specify this parameter, ElastiCache automatically chooses an appropriate time range.

Type: String

Required: No

TransitEncryptionEnabled

A flag that enables in-transit encryption when set to true. If you are enabling in-transit encryption for an existing cluster, you must also set `TransitEncryptionMode` to `preferred`.

Type: Boolean

Required: No

TransitEncryptionMode

A setting that allows you to migrate your clients to use in-transit encryption, with no downtime.

You must set `TransitEncryptionEnabled` to `true`, for your existing cluster, and set `TransitEncryptionMode` to `preferred` in the same request to allow both encrypted and unencrypted connections at the same time. Once you migrate all your Valkey or Redis OSS clients to use encrypted connections you can set the value to `required` to allow encrypted connections only.

Setting `TransitEncryptionMode` to `required` is a two-step process that requires you to first set the `TransitEncryptionMode` to `preferred`, after that you can set `TransitEncryptionMode` to `required`.

Type: String

Valid Values: `preferred | required`

Required: No

UserGroupIdsToAdd.member.N

The ID of the user group you are associating with the replication group.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

UserGroupIdsToRemove.member.N

The ID of the user group to disassociate from the replication group, meaning the users in the group no longer can access the replication group.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

Response Elements

The following element is returned by the service.

ReplicationGroup

Contains all of the attributes of a specific Valkey or Redis OSS replication group.

Type: ReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterNotFound

The requested cluster ID does not refer to an existing cluster.

HTTP Status Code: 404

CacheParameterGroupNotFound

The requested cache parameter group name does not refer to an existing cache parameter group.

HTTP Status Code: 404

CacheSecurityGroupNotFound

The requested cache security group name does not refer to an existing cache security group.

HTTP Status Code: 404

InsufficientCacheClusterCapacity

The requested cache node type is not available in the specified Availability Zone. For more information, see InsufficientCacheClusterCapacity in the ElastiCache User Guide.

HTTP Status Code: 400

InvalidCacheClusterState

The requested cluster is not in the `available` state.

HTTP Status Code: 400

InvalidCacheSecurityGroupState

The current state of the cache security group does not allow deletion.

HTTP Status Code: 400

InvalidKMSKeyFault

The KMS key supplied is not valid.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

InvalidUserGroupState

The user group is not in an active state.

HTTP Status Code: 400

InvalidVPCNetworkStateFault

The VPC network is in an invalid state.

HTTP Status Code: 400

NodeQuotaForClusterExceeded

The request cannot be processed because it would exceed the allowed number of cache nodes in a single cluster.

HTTP Status Code: 400

NodeQuotaForCustomerExceeded

The request cannot be processed because it would exceed the allowed number of cache nodes per customer.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

UserGroupNotFound

The user group was not found or does not exist

HTTP Status Code: 404

Examples

ModifyReplicationGroup

This example illustrates one usage of ModifyReplicationGroup.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=ModifyReplicationGroup
   &ApplyImmediately=false
   &ReplicationGroupId=my-repgroup
   &PrimaryClusterId=my-replica-1
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<ModifyReplicationGroupResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
   <ModifyReplicationGroupResult>
      <ReplicationGroup>
         <SnapshottingClusterId>my-redis-primary</SnapshottingClusterId>
         <MemberClusters>
            <ClusterId>my-redis-primary</ClusterId>
            <ClusterId>my-replica-1</ClusterId>
         </MemberClusters>
         <NodeGroups>
            <NodeGroup>
               <NodeGroupId>0001</NodeGroupId>
               <PrimaryEndpoint>
                  <Port>6379</Port>
                  <Address>my-repgroup.q68zge.ng.0001.use1devo.elmo-dev.amazonaws.com</
Address>
               </PrimaryEndpoint>
               <Status>available</Status>
               <NodeGroupMembers>
                  <NodeGroupMember>
                     <CacheClusterId>my-redis-primary</CacheClusterId>
                     <ReadEndpoint>
                        <Port>6379</Port>
                        <Address>my-redis-primary.q68zge.0001.use1devo.elmo-
dev.amazonaws.com</Address>
                     </ReadEndpoint>
                     <PreferredAvailabilityZone>us-west-2c</PreferredAvailabilityZone>
                     <CacheNodeId>0001</CacheNodeId>
```

```
                      <CurrentRole>primary</CurrentRole>
                   </NodeGroupMember>
                   <NodeGroupMember>
                      <CacheClusterId>my-replica-1</CacheClusterId>
                      <ReadEndpoint>
                         <Port>6379</Port>
                         <Address>my-replica-1.q68zge.0001.use1devo.elmo-
dev.amazonaws.com</Address>
                      </ReadEndpoint>
                      <PreferredAvailabilityZone>us-west-2b</PreferredAvailabilityZone>
                      <CacheNodeId>0001</CacheNodeId>
                      <CurrentRole>replica</CurrentRole>
                   </NodeGroupMember>
                </NodeGroupMembers>
             </NodeGroup>
          </NodeGroups>
          <ReplicationGroupId>my-repgroup</ReplicationGroupId>
          <Status>available</Status>
          <PendingModifiedValues>
             <PrimaryClusterId>my-replica-1</PrimaryClusterId>
          </PendingModifiedValues>
          <Description>My replication group</Description>
       </ReplicationGroup>
    </ModifyReplicationGroupResult>
    <ResponseMetadata>
       <RequestId>6fd0aad6-b9d7-11e3-8a16-7978bb24ffdf</RequestId>
    </ResponseMetadata>
</ModifyReplicationGroupResponse>
```

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

ModifyReplicationGroupShardConfiguration

Modifies a replication group's shards (node groups) by allowing you to add shards, remove shards, or rebalance the keyspaces among existing shards.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ApplyImmediately

Indicates that the shard reconfiguration process begins immediately. At present, the only permitted value for this parameter is `true`.

Value: true

Type: Boolean

Required: Yes

NodeGroupCount

The number of node groups (shards) that results from the modification of the shard configuration.

Type: Integer

Required: Yes

ReplicationGroupId

The name of the Valkey or Redis OSS (cluster mode enabled) cluster (replication group) on which the shards are to be configured.

Type: String

Required: Yes

NodeGroupsToRemove.NodeGroupToRemove.N

If the value of `NodeGroupCount` is less than the current number of node groups (shards), then either `NodeGroupsToRemove` or `NodeGroupsToRetain` is required. `NodeGroupsToRemove` is a list of `NodeGroupIds` to remove from the cluster.

ElastiCache will attempt to remove all node groups listed by `NodeGroupsToRemove` from the cluster.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 4.

Pattern: `\d+`

Required: No

NodeGroupsToRetain.NodeGroupToRetain.N

If the value of `NodeGroupCount` is less than the current number of node groups (shards), then either `NodeGroupsToRemove` or `NodeGroupsToRetain` is required. `NodeGroupsToRetain` is a list of `NodeGroupIds` to retain in the cluster.

ElastiCache will attempt to remove all node groups except those listed by `NodeGroupsToRetain` from the cluster.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 4.

Pattern: `\d+`

Required: No

ReshardingConfiguration.ReshardingConfiguration.N

Specifies the preferred availability zones for each node group in the cluster. If the value of `NodeGroupCount` is greater than the current number of node groups (shards), you can use this parameter to specify the preferred availability zones of the cluster's shards. If you omit this parameter ElastiCache selects availability zones for you.

You can specify this parameter only if the value of `NodeGroupCount` is greater than the current number of node groups (shards).

Type: Array of ReshardingConfiguration objects

Required: No

Response Elements

The following element is returned by the service.

ReplicationGroup

Contains all of the attributes of a specific Valkey or Redis OSS replication group.

Type: ReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

InsufficientCacheClusterCapacity

The requested cache node type is not available in the specified Availability Zone. For more information, see InsufficientCacheClusterCapacity in the ElastiCache User Guide.

HTTP Status Code: 400

InvalidCacheClusterState

The requested cluster is not in the `available` state.

HTTP Status Code: 400

InvalidKMSKeyFault

The KMS key supplied is not valid.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

InvalidVPCNetworkStateFault

The VPC network is in an invalid state.

HTTP Status Code: 400

NodeGroupsPerReplicationGroupQuotaExceeded

The request cannot be processed because it would exceed the maximum allowed number of node groups (shards) in a single replication group. The default maximum is 90

HTTP Status Code: 400

NodeQuotaForCustomerExceeded

The request cannot be processed because it would exceed the allowed number of cache nodes per customer.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

Examples

Add shards

The following example adds shards to the replication group `my-cluster` so that at the completion of the call there are 4 shards. The availability zones for the nodes in the shards are specified by the parameter `ReshardingConfiguration.ReshardingConfiguration.N.PreferredAvailabilityZones.Ava` If there are already 4 shards in this replication group, the call fails.

```
https://elasticache.us-east-2.amazonaws.com/
    ?Action=ModifyReplicationGroupShardConfiguration
```

```
    &ApplyImmediately=true
    &NodeGroupCount=4
    &ReplicationGroupId=my-cluster

 &ReshardingConfiguration.ReshardingConfiguration.1.PreferredAvailabilityZones.AvailabilityZone
east-2a

 &ReshardingConfiguration.ReshardingConfiguration.1.PreferredAvailabilityZones.AvailabilityZone
east-2c

 &ReshardingConfiguration.ReshardingConfiguration.2.PreferredAvailabilityZones.AvailabilityZone
east-2b

 &ReshardingConfiguration.ReshardingConfiguration.2.PreferredAvailabilityZones.AvailabilityZone
east-2a

 &ReshardingConfiguration.ReshardingConfiguration.3.PreferredAvailabilityZones.AvailabilityZone
east-2c

 &ReshardingConfiguration.ReshardingConfiguration.3.PreferredAvailabilityZones.AvailabilityZone
east-2d

 &ReshardingConfiguration.ReshardingConfiguration.4.PreferredAvailabilityZones.AvailabilityZone
east-2d

 &ReshardingConfiguration.ReshardingConfiguration.4.PreferredAvailabilityZones.AvailabilityZone
east-2c
    &Version=2015-02-02
    &SignatureVersion=4
    &SignatureMethod=HmacSHA256
    &Timestamp=20171002T192317Z
    &X-Amz-Credential=<credential>
```

Remove shards

The following example removes two shards from the replication group `my-cluster`. leaving 2 shards. When removing shards, the parameter `NodeGroupsToRemove.NodeGroupToRemove` is required.

You cannot remove the last shard leaving zero shards.

```
https://elasticache.us-east-2.amazonaws.com/
    ?Action=ModifyReplicationGroupShardConfiguration
```

```
    &ApplyImmediately=true
    &NodeGroupCount=2
    &ReplicationGroupId=my-cluster
    &NodeGroupsToRemove.NodeGroupToRemove.1=0002
    &NodeGroupsToRemove.NodeGroupToRemove.2=0003
    &Version=2015-02-02
    &SignatureVersion=4
    &SignatureMethod=HmacSHA256
    &Timestamp=20171002T192317Z
    &X-Amz-Credential=<credential>
```

Rebalance shards

The following rebalances the keyspaces among the existing shards in the replication group `my-cluster`. The value specified by `NodeGroupCount` must be the existing number of shards. If the keyspaces are already balanced the call fails.

```
https://elasticache.us-east-2.amazonaws.com/
    ?Action=ModifyReplicationGroupShardConfiguration
    &ApplyImmediately=true
    &NodeGroupCount=4
    &ReplicationGroupId=my-cluster
    &Version=2015-02-02
    &SignatureVersion=4
    &SignatureMethod=HmacSHA256
    &Timestamp=20171002T192317Z
    &X-Amz-Credential=<credential>
```

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

ModifyServerlessCache

This API modifies the attributes of a serverless cache.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ServerlessCacheName

User-provided identifier for the serverless cache to be modified.

Type: String

Required: Yes

CacheUsageLimits

Modify the cache usage limit for the serverless cache.

Type: CacheUsageLimits object

Required: No

DailySnapshotTime

The daily time during which Elasticache begins taking a daily snapshot of the serverless cache. Available for Valkey, Redis OSS and Serverless Memcached only. The default is NULL, i.e. the existing snapshot time configured for the cluster is not removed.

Type: String

Required: No

Description

User provided description for the serverless cache. Default = NULL, i.e. the existing description is not removed/modified. The description has a maximum length of 255 characters.

Type: String

Required: No

Engine

Modifies the engine listed in a serverless cache request. The options are valkey, memcached or redis.

Type: String

Required: No

MajorEngineVersion

Modifies the engine vesion listed in a serverless cache request.

Type: String

Required: No

RemoveUserGroup

The identifier of the UserGroup to be removed from association with the Valkey and Redis OSS serverless cache. Available for Valkey and Redis OSS only. Default is NULL.

Type: Boolean

Required: No

SecurityGroupIds.SecurityGroupId.N

The new list of VPC security groups to be associated with the serverless cache. Populating this list means the current VPC security groups will be removed. This security group is used to authorize traffic access for the VPC end-point (private-link). Default = NULL - the existing list of VPC security groups is not removed.

Type: Array of strings

Required: No

SnapshotRetentionLimit

The number of days for which Elasticache retains automatic snapshots before deleting them. Available for Valkey, Redis OSS and Serverless Memcached only. Default = NULL, i.e. the existing snapshot-retention-limit will not be removed or modified. The maximum value allowed is 35 days.

Type: Integer

Required: No

UserGroupId

The identifier of the UserGroup to be associated with the serverless cache. Available for Valkey and Redis OSS only. Default is NULL - the existing UserGroup is not removed.

Type: String

Required: No

Response Elements

The following element is returned by the service.

ServerlessCache

The response for the attempt to modify the serverless cache.

Type: ServerlessCache object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidCredentials

You must enter valid credentials.

HTTP Status Code: 408

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidServerlessCacheStateFault

The account for these credentials is not currently active.

HTTP Status Code: 400

InvalidUserGroupState

The user group is not in an active state.

HTTP Status Code: 400

ServerlessCacheNotFoundFault

The serverless cache was not found or does not exist.

HTTP Status Code: 404

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

UserGroupNotFound

The user group was not found or does not exist

HTTP Status Code: 404

Examples

ModifyServerlessCache

This example illustrates one usage of ModifyServerlessCache.

Sample Request

```
{
    "description": "Modify the daily snapshot time of cache",
    "input": {
        "ServerlessCacheName": "my-serverless-cache",
        "CacheUsageLimits": {
            "DataStorage" : {
                "Maximum" : 10,
                "Unit" : "GB"
            },
            "ECPUPerSecond" : {
                "Maximum" : 50000
            }
        },
        "DailySnapshotTime": "11:00"
    },
    "output": {
```

```
        "ServerlessCache": {
            "ServerlessCacheName": "my-serverless-cache",
            "Description": "A serverless cache.",
            "Status": "available",
            "Engine": "redis",
            "MajorEngineVersion": "7",
            "FullEngineVersion": "7.0",
            "SubnetIds": [
                "subnet-xxx8c982",
                "subnet-xxx382f3",
                "subnet-xxxb3e7c0"
            ],
            "CacheUsageLimits": {
                "DataStorage" : {
                    "Maximum" : 10,
                    "Unit" : "GB"
                },
                "ECPUPerSecond" : {
                    "Maximum" : 50000
                }
            },
            "SecurityGroupIds": [
                "sg-xxx0c9af"
            ],
            "Endpoint": {
                "Address": "my-serverless-cache-
xxxxxx.serverless.use1qa.cache.amazonaws.com",
                "Port": 6379
            },
            "ARN": "arn:aws:elasticache:us-east-1:222222222222:serverlesscache:my-
serverless-cache",
            "SnapshotRetentionLimit": 10,
            "DailySnapshotTime": "11:00"
        }
    }
}
```

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

ModifyUser

Changes user password(s) and/or access string.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

UserId

The ID of the user.

Type: String

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: Yes

AccessString

Access permissions string used for this user.

Type: String

Pattern: `.*\S.*`

Required: No

AppendAccessString

Adds additional user permissions to the access string.

Type: String

Pattern: `.*\S.*`

Required: No

AuthenticationMode

Specifies how to authenticate the user.

Type: AuthenticationMode object

Required: No

Engine

Modifies the engine listed for a user. The options are valkey or redis.

Type: String

Pattern: `[a-zA-Z]*`

Required: No

NoPasswordRequired

Indicates no password is required for the user.

Type: Boolean

Required: No

Passwords.member.N

The passwords belonging to the user. You are allowed up to two.

Type: Array of strings

Array Members: Minimum number of 1 item.

Required: No

Response Elements

The following elements are returned by the service.

AccessString

Access permissions string used for this user.

Type: String

ARN

The Amazon Resource Name (ARN) of the user.

Type: String

Authentication

Denotes whether the user requires a password to authenticate.

Type: Authentication object

Engine

The options are valkey or redis.

Type: String

Pattern: `[a-zA-Z]*`

MinimumEngineVersion

The minimum engine version required, which is Redis OSS 6.0

Type: String

Status

Indicates the user status. Can be "active", "modifying" or "deleting".

Type: String

UserGroupIds.member.N

Returns a list of the user group IDs the user belongs to.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

UserId

The ID of the user.

Type: String

UserName

The username of the user.

Type: String

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidUserState

The user is not in active state.

HTTP Status Code: 400

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

UserNotFound

The user does not exist or could not be found.

HTTP Status Code: 404

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

ModifyUserGroup

Changes the list of users that belong to the user group.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

UserGroupId

The ID of the user group.

Type: String

Required: Yes

Engine

Modifies the engine listed in a user group. The options are valkey or redis.

Type: String

Pattern: `[a-zA-Z]*`

Required: No

UserIdsToAdd.member.N

The list of user IDs to add to the user group.

Type: Array of strings

Array Members: Minimum number of 1 item.

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

UserIdsToRemove.member.N

The list of user IDs to remove from the user group.

Type: Array of strings

Array Members: Minimum number of 1 item.

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

Response Elements

The following elements are returned by the service.

ARN

The Amazon Resource Name (ARN) of the user group.

Type: String

Engine

The options are valkey or redis.

Type: String

Pattern: `[a-zA-Z]*`

MinimumEngineVersion

The minimum engine version required, which is Redis OSS 6.0

Type: String

PendingChanges

A list of updates being applied to the user group.

Type: UserGroupPendingChanges object

ReplicationGroups.member.N

A list of replication groups that the user group can access.

Type: Array of strings

ServerlessCaches.member.N

Indicates which serverless caches the specified user group is associated with. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: Array of strings

Status

Indicates user group status. Can be "creating", "active", "modifying", "deleting".

Type: String

UserGroupId

The ID of the user group.

Type: String

UserIds.member.N

The list of user IDs that belong to the user group.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Errors

For information about the errors that are common to all actions, see Common Errors.

DefaultUserRequired

You must add default user to a user group.

HTTP Status Code: 400

DuplicateUserName

A user with this username already exists.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidUserGroupState

The user group is not in an active state.

HTTP Status Code: 400

ServiceLinkedRoleNotFoundFault

The specified service linked role (SLR) was not found.

HTTP Status Code: 400

UserGroupNotFound

The user group was not found or does not exist

HTTP Status Code: 404

UserNotFound

The user does not exist or could not be found.

HTTP Status Code: 404

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

PurchaseReservedCacheNodesOffering

Allows you to purchase a reserved cache node offering. Reserved nodes are not eligible for cancellation and are non-refundable. For more information, see Managing Costs with Reserved Nodes.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ReservedCacheNodesOfferingId

The ID of the reserved cache node offering to purchase.

Example: `438012d3-4052-4cc7-b2e3-8d3372e0e706`

Type: String

Required: Yes

CacheNodeCount

The number of cache node instances to reserve.

Default: 1

Type: Integer

Required: No

ReservedCacheNodeId

A customer-specified identifier to track this reservation.

> **Note**
>
> The Reserved Cache Node ID is an unique customer-specified identifier to track this reservation. If this parameter is not specified, ElastiCache automatically generates an identifier for the reservation.

Example: myreservationID

Type: String

Required: No

Tags.Tag.N

A list of tags to be added to this resource. A tag is a key-value pair. A tag key must be accompanied by a tag value, although null is accepted.

Type: Array of Tag objects

Required: No

Response Elements

The following element is returned by the service.

ReservedCacheNode

Represents the output of a `PurchaseReservedCacheNodesOffering` operation.

Type: ReservedCacheNode object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

ReservedCacheNodeAlreadyExists

You already have a reservation with the given identifier.

HTTP Status Code: 404

ReservedCacheNodeQuotaExceeded

The request cannot be processed because it would exceed the user's cache node quota.

HTTP Status Code: 400

ReservedCacheNodesOfferingNotFound

The requested cache node offering does not exist.

HTTP Status Code: 404

TagQuotaPerResourceExceeded

The request cannot be processed because it would cause the resource to have more than the allowed number of tags. The maximum number of tags permitted on a resource is 50.

HTTP Status Code: 400

Examples

PurchaseReservedCacheNodesOffering

This example illustrates one usage of PurchaseReservedCacheNodesOffering.

Sample Request

```
https://elasticache.amazonaws.com/
   ?Action=PurchaseReservedCacheNodesOffering
   &ReservedCacheNodeId=myreservationID
   &ReservedCacheNodesOfferingId=438012d3-4052-4cc7-b2e3-8d3372e0e706
   &CacheNodeCount=1
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<PurchaseReservedCacheNodesOfferingResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
  <PurchaseReservedCacheNodesOfferingResult>
```

```
    <ReservedCacheNode>
      <OfferingType>Medium Utilization</OfferingType>
      <RecurringCharges/>
      <ProductDescription>memcached</ProductDescription>
      <ReservedCacheNodesOfferingId> 438012d3-4052-4cc7-b2e3-8d3372e0e706</
ReservedCacheNodesOfferingId>
      <State>payment-pending</State>
      <ReservedCacheNodeId>myreservationID</ReservedCacheNodeId>
      <CacheNodeCount>10</CacheNodeCount>
      <StartTime>2015-02-02T23:24:56.577Z</StartTime>
      <Duration>31536000</Duration>
      <FixedPrice>123.0</FixedPrice>
      <UsagePrice>0.123</UsagePrice>
      <CacheNodeType>cache.m1.small</CacheNodeType>
    </ReservedCacheNode>
  </PurchaseReservedCacheNodesOfferingResult>
  <ResponseMetadata>
    <RequestId>7f099901-29cf-11e1-bd06-6fe008f046c3</RequestId>
  </ResponseMetadata>
</PurchaseReservedCacheNodesOfferingResponse>
```

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

RebalanceSlotsInGlobalReplicationGroup

Redistribute slots to ensure uniform distribution across existing shards in the cluster.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ApplyImmediately

If `True`, redistribution is applied immediately.

Type: Boolean

Required: Yes

GlobalReplicationGroupId

The name of the Global datastore

Type: String

Required: Yes

Response Elements

The following element is returned by the service.

GlobalReplicationGroup

Consists of a primary cluster that accepts writes and an associated secondary cluster that resides in a different Amazon region. The secondary cluster accepts only reads. The primary cluster automatically replicates updates to the secondary cluster.

- The **GlobalReplicationGroupIdSuffix** represents the name of the Global datastore, which is what you use to associate a secondary cluster.

Type: GlobalReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

GlobalReplicationGroupNotFoundFault

The Global datastore does not exist

HTTP Status Code: 404

InvalidGlobalReplicationGroupState

The Global datastore is not available or in primary-only state.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

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

RebootCacheCluster

Reboots some, or all, of the cache nodes within a provisioned cluster. This operation applies any modified cache parameter groups to the cluster. The reboot operation takes place as soon as possible, and results in a momentary outage to the cluster. During the reboot, the cluster status is set to REBOOTING.

The reboot causes the contents of the cache (for each cache node being rebooted) to be lost.

When the reboot is complete, a cluster event is created.

Rebooting a cluster is currently supported on Memcached, Valkey and Redis OSS (cluster mode disabled) clusters. Rebooting is not supported on Valkey or Redis OSS (cluster mode enabled) clusters.

If you make changes to parameters that require a Valkey or Redis OSS (cluster mode enabled) cluster reboot for the changes to be applied, see Rebooting a Cluster for an alternate process.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheClusterId

The cluster identifier. This parameter is stored as a lowercase string.

Type: String

Required: Yes

CacheNodeIdsToReboot.CacheNodeId.N

A list of cache node IDs to reboot. A node ID is a numeric identifier (0001, 0002, etc.). To reboot an entire cluster, specify all of the cache node IDs.

Type: Array of strings

Required: Yes

Response Elements

The following element is returned by the service.

CacheCluster

Contains all of the attributes of a specific cluster.

Type: CacheCluster object

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterNotFound

The requested cluster ID does not refer to an existing cluster.

HTTP Status Code: 404

InvalidCacheClusterState

The requested cluster is not in the `available` state.

HTTP Status Code: 400

Examples

RebootCacheCluster

This example illustrates one usage of RebootCacheCluster.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=RebootCacheCluster
   &CacheClusterId=mycache
   &CacheNodeIdsToReboot.CacheNodeId.1=0001
   &CacheNodeIdsToReboot.CacheNodeId.2=0002
   &CacheNodeIdsToReboot.CacheNodeId.3=0003
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response


```
<RebootCacheClusterResponse xmlns="http://elasticache.amazonaws.com/doc/2015-02-02/">
   <RebootCacheClusterResult>
      <CacheCluster>
         <CacheClusterStatus>rebooting cluster nodes</CacheClusterStatus>
         <CacheParameterGroup>
            <CacheParameterGroupName>default.memcached1.4</CacheParameterGroupName>
            <ParameterApplyStatus>in-sync</ParameterApplyStatus>
            <CacheNodeIdsToReboot />
         </CacheParameterGroup>
         <CacheClusterId>mycache</CacheClusterId>
         <ConfigurationEndpoint>
            <Port>11211</Port>
            <Address>mycache.q68zge.cfg.use1devo.elmo-dev.amazonaws.com</Address>
         </ConfigurationEndpoint>
         <CacheNodeType>cache.m1.small</CacheNodeType>
         <Engine>memcached</Engine>
         <PendingModifiedValues />
         <PreferredAvailabilityZone>us-west-2b</PreferredAvailabilityZone>
         <CacheClusterCreateTime>2015-02-02T19:04:12.812Z</CacheClusterCreateTime>
         <EngineVersion>1.4.17</EngineVersion>
         <AutoMinorVersionUpgrade>true</AutoMinorVersionUpgrade>
         <PreferredMaintenanceWindow>wed:09:00-wed:10:00</PreferredMaintenanceWindow>
         <ClientDownloadLandingPage>https://console.aws.amazon.com/elasticache/
home#client-download:</ClientDownloadLandingPage>
         <CacheSecurityGroups>
            <CacheSecurityGroup>
               <CacheSecurityGroupName>default</CacheSecurityGroupName>
               <Status>active</Status>
            </CacheSecurityGroup>
         </CacheSecurityGroups>
         <NumCacheNodes>3</NumCacheNodes>
      </CacheCluster>
   </RebootCacheClusterResult>
   <ResponseMetadata>
      <RequestId>cf7e6fc4-b9d1-11e3-8a16-7978bb24ffdf</RequestId>
   </ResponseMetadata>
</RebootCacheClusterResponse>
```

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

RemoveTagsFromResource

Removes the tags identified by the `TagKeys` list from the named resource. A tag is a key-value pair where the key and value are case-sensitive. You can use tags to categorize and track all your ElastiCache resources, with the exception of global replication group. When you add or remove tags on replication groups, those actions will be replicated to all nodes in the replication group. For more information, see Resource-level permissions.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

ResourceName

The Amazon Resource Name (ARN) of the resource from which you want the tags removed, for example `arn:aws:elasticache:us-west-2:0123456789:cluster:myCluster` or `arn:aws:elasticache:us-west-2:0123456789:snapshot:mySnapshot`.

For more information about ARNs, see Amazon Resource Names (ARNs) and Amazon Service Namespaces.

Type: String

Required: Yes

TagKeys.member.N

A list of `TagKeys` identifying the tags you want removed from the named resource.

Type: Array of strings

Required: Yes

Response Elements

The following element is returned by the service.

TagList.Tag.N

A list of tags as key-value pairs.

Type: Array of Tag objects

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheClusterNotFound

The requested cluster ID does not refer to an existing cluster.

HTTP Status Code: 404

CacheParameterGroupNotFound

The requested cache parameter group name does not refer to an existing cache parameter group.

HTTP Status Code: 404

CacheSecurityGroupNotFound

The requested cache security group name does not refer to an existing cache security group.

HTTP Status Code: 404

CacheSubnetGroupNotFoundFault

The requested cache subnet group name does not refer to an existing cache subnet group.

HTTP Status Code: 400

InvalidARN

The requested Amazon Resource Name (ARN) does not refer to an existing resource.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

InvalidServerlessCacheSnapshotStateFault

The state of the serverless cache snapshot was not received. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 400

InvalidServerlessCacheStateFault

The account for these credentials is not currently active.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

ReservedCacheNodeNotFound

The requested reserved cache node was not found.

HTTP Status Code: 404

ServerlessCacheNotFoundFault

The serverless cache was not found or does not exist.

HTTP Status Code: 404

ServerlessCacheSnapshotNotFoundFault

This serverless cache snapshot could not be found or does not exist. Available for Valkey, Redis OSS and Serverless Memcached only.

HTTP Status Code: 404

SnapshotNotFoundFault

The requested snapshot name does not refer to an existing snapshot.

HTTP Status Code: 404

TagNotFound

The requested tag was not found on this resource.

HTTP Status Code: 404

UserGroupNotFound

The user group was not found or does not exist

HTTP Status Code: 404

UserNotFound

The user does not exist or could not be found.

HTTP Status Code: 404

Examples

RemoveTagsFromResource

This example illustrates one usage of RemoveTagsFromResource.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=RemoveTagsFromResource
   &ResourceName=arn:aws:elasticache:us-west-2:0123456789:cluster:myCluster
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &TagKeys.TagKey.1=service
   &TagKeys.TagKey.2=organization
   &Version=2015-02-02
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

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

ResetCacheParameterGroup

Modifies the parameters of a cache parameter group to the engine or system default value. You can reset specific parameters by submitting a list of parameter names. To reset the entire cache parameter group, specify the `ResetAllParameters` and `CacheParameterGroupName` parameters.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheParameterGroupName

The name of the cache parameter group to reset.

Type: String

Required: Yes

ParameterNameValues.ParameterNameValue.N

An array of parameter names to reset to their default values. If `ResetAllParameters` is `true`, do not use `ParameterNameValues`. If `ResetAllParameters` is `false`, you must specify the name of at least one parameter to reset.

Type: Array of ParameterNameValue objects

Required: No

ResetAllParameters

If `true`, all parameters in the cache parameter group are reset to their default values. If `false`, only the parameters listed by `ParameterNameValues` are reset to their default values.

Valid values: `true` | `false`

Type: Boolean

Required: No

Response Elements

The following element is returned by the service.

CacheParameterGroupName

The name of the cache parameter group.

Type: String

Errors

For information about the errors that are common to all actions, see Common Errors.

CacheParameterGroupNotFound

The requested cache parameter group name does not refer to an existing cache parameter group.

HTTP Status Code: 404

InvalidCacheParameterGroupState

The current state of the cache parameter group does not allow the requested operation to occur.

HTTP Status Code: 400

InvalidGlobalReplicationGroupState

The Global datastore is not available or in primary-only state.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

Examples

ResetCacheParameterGroup

This example illustrates one usage of ResetCacheParameterGroup.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=ResetCacheParameterGroup
   &ResetAllParameters=true
   &CacheParameterGroupName=mycacheparametergroup1
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<ResetCacheParameterGroupResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
   <ResetCacheParameterGroupResult>
      <CacheParameterGroupName>mycacheparametergroup1</CacheParameterGroupName>
   </ResetCacheParameterGroupResult>
   <ResponseMetadata>
      <RequestId>cb7cc855-b9d2-11e3-8a16-7978bb24ffdf</RequestId>
   </ResponseMetadata>
</ResetCacheParameterGroupResponse>
```

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

RevokeCacheSecurityGroupIngress

Revokes ingress from a cache security group. Use this operation to disallow access from an Amazon EC2 security group that had been previously authorized.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CacheSecurityGroupName

The name of the cache security group to revoke ingress from.

Type: String

Required: Yes

EC2SecurityGroupName

The name of the Amazon EC2 security group to revoke access from.

Type: String

Required: Yes

EC2SecurityGroupOwnerId

The Amazon account number of the Amazon EC2 security group owner. Note that this is not the same thing as an Amazon access key ID - you must provide a valid Amazon account number for this parameter.

Type: String

Required: Yes

Response Elements

The following element is returned by the service.

CacheSecurityGroup

Represents the output of one of the following operations:

- `AuthorizeCacheSecurityGroupIngress`

- `CreateCacheSecurityGroup`
- `RevokeCacheSecurityGroupIngress`

Type: CacheSecurityGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

AuthorizationNotFound

The specified Amazon EC2 security group is not authorized for the specified cache security group.

HTTP Status Code: 404

CacheSecurityGroupNotFound

The requested cache security group name does not refer to an existing cache security group.

HTTP Status Code: 404

InvalidCacheSecurityGroupState

The current state of the cache security group does not allow deletion.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

Examples

RevokeCacheSecurityGroupIngress

This example illustrates one usage of RevokeCacheSecurityGroupIngress.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=RevokeCacheSecurityGroupIngress
   &EC2SecurityGroupName=default
   &CacheSecurityGroupName=mygroup
   &EC2SecurityGroupOwnerId=1234-5678-1234
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20150202T192317Z
   &X-Amz-Credential=<credential>
```

Sample Response

```
<RevokeCacheSecurityGroupIngressResponse xmlns="http://elasticache.amazonaws.com/
doc/2015-02-02/">
    <RevokeCacheSecurityGroupIngressResult>
        <CacheSecurityGroup>
            <EC2SecurityGroups>
                <EC2SecurityGroup>
                    <Status>revoking</Status>
                    <EC2SecurityGroupName>default</EC2SecurityGroupName>
                    <EC2SecurityGroupOwnerId>123456781234</EC2SecurityGroupOwnerId>
                </EC2SecurityGroup>
            </EC2SecurityGroups>
            <CacheSecurityGroupName>mygroup</CacheSecurityGroupName>
            <OwnerId>123456789012</OwnerId>
            <Description>My security group</Description>
        </CacheSecurityGroup>
    </RevokeCacheSecurityGroupIngressResult>
    <ResponseMetadata>
        <RequestId>02ae3699-3650-11e0-a564-8f11342c56b0</RequestId>
    </ResponseMetadata>
</RevokeCacheSecurityGroupIngressResponse>
```

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

StartMigration

Start the migration of data.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CustomerNodeEndpointList.member.N

List of endpoints from which data should be migrated. For Valkey or Redis OSS (cluster mode disabled), the list should have only one element.

Type: Array of CustomerNodeEndpoint objects

Required: Yes

ReplicationGroupId

The ID of the replication group to which data should be migrated.

Type: String

Required: Yes

Response Elements

The following element is returned by the service.

ReplicationGroup

Contains all of the attributes of a specific Valkey or Redis OSS replication group.

Type: ReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

ReplicationGroupAlreadyUnderMigrationFault

The targeted replication group is not available.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

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

TestFailover

Represents the input of a `TestFailover` operation which tests automatic failover on a specified node group (called shard in the console) in a replication group (called cluster in the console).

This API is designed for testing the behavior of your application in case of ElastiCache failover. It is not designed to be an operational tool for initiating a failover to overcome a problem you may have with the cluster. Moreover, in certain conditions such as large-scale operational events, Amazon may block this API.

Note the following

- A customer can use this operation to test automatic failover on up to 15 shards (called node groups in the ElastiCache API and Amazon CLI) in any rolling 24-hour period.
- If calling this operation on shards in different clusters (called replication groups in the API and CLI), the calls can be made concurrently.
- If calling this operation multiple times on different shards in the same Valkey or Redis OSS (cluster mode enabled) replication group, the first node replacement must complete before a subsequent call can be made.
- To determine whether the node replacement is complete you can check Events using the Amazon ElastiCache console, the Amazon CLI, or the ElastiCache API. Look for the following automatic failover related events, listed here in order of occurrance:
  1. Replication group message: `Test Failover API called for node group <node-group-id>`
  2. Cache cluster message: `Failover from primary node <primary-node-id> to replica node <node-id> completed`
  3. Replication group message: `Failover from primary node <primary-node-id> to replica node <node-id> completed`
  4. Cache cluster message: `Recovering cache nodes <node-id>`
  5. Cache cluster message: `Finished recovery for cache nodes <node-id>`

  For more information see:

  - Viewing ElastiCache Events in the *ElastiCache User Guide*
  - DescribeEvents in the ElastiCache API Reference

Also see, Testing Multi-AZ in the *ElastiCache User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

NodeGroupId

The name of the node group (called shard in the console) in this replication group on which automatic failover is to be tested. You may test automatic failover on up to 15 node groups in any rolling 24-hour period.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 4.

Pattern: \d+

Required: Yes

ReplicationGroupId

The name of the replication group (console: cluster) whose automatic failover is being tested by this operation.

Type: String

Required: Yes

Response Elements

The following element is returned by the service.

ReplicationGroup

Contains all of the attributes of a specific Valkey or Redis OSS replication group.

Type: ReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

APICallRateForCustomerExceeded

The customer has exceeded the allowed rate of API calls.

HTTP Status Code: 400

InvalidCacheClusterState

The requested cluster is not in the `available` state.

HTTP Status Code: 400

InvalidKMSKeyFault

The KMS key supplied is not valid.

HTTP Status Code: 400

InvalidParameterCombination

Two or more incompatible parameters were specified.

HTTP Status Code: 400

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

NodeGroupNotFoundFault

The node group specified by the `NodeGroupId` parameter could not be found. Please verify that the node group exists and that you spelled the `NodeGroupId` value correctly.

HTTP Status Code: 404

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

TestFailoverNotAvailableFault

The `TestFailover` action is not available.

HTTP Status Code: 400

Examples

Example

The following example tests automatic failover on the Valkey or Redis OSS (cluster mode disabled) replication group (console: cluster) `redis00`. Since there is only one node group in Valkey or Redis OSS (cluster mode disabled) clusters, the *NodeGroupId* will always be `<cluster-name>-0001`.

Sample Request

```
https://elasticache.us-west-2.amazonaws.com/
   ?Action=TestFailover
   &NodeGroupId=redis00-0001
   &ReplicationGroupId=redis00
   &Version=2015-02-02
   &SignatureVersion=4
   &SignatureMethod=HmacSHA256
   &Timestamp=20170418T192317Z
   &X-Amz-Credential=<credential>
```

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

TestMigration

Async API to test connection between source and target replication group.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

CustomerNodeEndpointList.member.N

List of endpoints from which data should be migrated. List should have only one element.

Type: Array of CustomerNodeEndpoint objects

Required: Yes

ReplicationGroupId

The ID of the replication group to which data is to be migrated.

Type: String

Required: Yes

Response Elements

The following element is returned by the service.

ReplicationGroup

Contains all of the attributes of a specific Valkey or Redis OSS replication group.

Type: ReplicationGroup object

Errors

For information about the errors that are common to all actions, see Common Errors.

InvalidParameterValue

The value for a parameter is invalid.

HTTP Status Code: 400

InvalidReplicationGroupState

The requested replication group is not in the `available` state.

HTTP Status Code: 400

ReplicationGroupAlreadyUnderMigrationFault

The targeted replication group is not available.

HTTP Status Code: 400

ReplicationGroupNotFoundFault

The specified replication group does not exist.

HTTP Status Code: 404

See Also

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

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

Data Types

The Amazon ElastiCache 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:

- Authentication
- AuthenticationMode
- AvailabilityZone
- CacheCluster
- CacheEngineVersion
- CacheNode
- CacheNodeTypeSpecificParameter
- CacheNodeTypeSpecificValue
- CacheNodeUpdateStatus
- CacheParameterGroup
- CacheParameterGroupStatus
- CacheSecurityGroup
- CacheSecurityGroupMembership
- CacheSubnetGroup
- CacheUsageLimits
- CloudWatchLogsDestinationDetails
- ConfigureShard
- CustomerNodeEndpoint
- DataStorage
- DestinationDetails

- EC2SecurityGroup
- ECPUPerSecond
- Endpoint
- EngineDefaults
- Event
- Filter
- GlobalNodeGroup
- GlobalReplicationGroup
- GlobalReplicationGroupInfo
- GlobalReplicationGroupMember
- KinesisFirehoseDestinationDetails
- LogDeliveryConfiguration
- LogDeliveryConfigurationRequest
- NodeGroup
- NodeGroupConfiguration
- NodeGroupMember
- NodeGroupMemberUpdateStatus
- NodeGroupUpdateStatus
- NodeSnapshot
- NotificationConfiguration
- Parameter
- ParameterNameValue
- PendingLogDeliveryConfiguration
- PendingModifiedValues
- ProcessedUpdateAction
- RecurringCharge
- RegionalConfiguration
- ReplicationGroup
- ReplicationGroupPendingModifiedValues
- ReservedCacheNode

- ReservedCacheNodesOffering
- ReshardingConfiguration
- ReshardingStatus
- ScaleConfig
- SecurityGroupMembership
- ServerlessCache
- ServerlessCacheConfiguration
- ServerlessCacheSnapshot
- ServiceUpdate
- SlotMigration
- Snapshot
- Subnet
- SubnetOutpost
- Tag
- TimeRangeFilter
- UnprocessedUpdateAction
- UpdateAction
- User
- UserGroup
- UserGroupPendingChanges
- UserGroupsUpdateStatus

Authentication

Indicates whether the user requires a password to authenticate.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

PasswordCount

The number of passwords belonging to the user. The maximum is two.

Type: Integer

Required: No

Type

Indicates whether the user requires a password to authenticate.

Type: String

Valid Values: `password | no-password | iam`

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

AuthenticationMode

Specifies the authentication mode to use.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

Passwords.member.N

Specifies the passwords to use for authentication if `Type` is set to `password`.

Type: Array of strings

Array Members: Minimum number of 1 item.

Required: No

Type

Specifies the authentication type. Possible options are IAM authentication, password and no password.

Type: String

Valid Values: `password | no-password-required | iam`

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

AvailabilityZone

Describes an Availability Zone in which the cluster is launched.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

Name

The name of the 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

CacheCluster

Contains all of the attributes of a specific cluster.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

ARN

The ARN (Amazon Resource Name) of the cache cluster.

Type: String

Required: No

AtRestEncryptionEnabled

A flag that enables encryption at-rest when set to `true`.

You cannot modify the value of `AtRestEncryptionEnabled` after the cluster is created. To enable at-rest encryption on a cluster you must set `AtRestEncryptionEnabled` to `true` when you create a cluster.

Required: Only available when creating a replication group in an Amazon VPC using Redis OSS version `3.2.6`, `4.x` or later.

Default: `false`

Type: Boolean

Required: No

AuthTokenEnabled

A flag that enables using an `AuthToken` (password) when issuing Valkey or Redis OSS commands.

Default: `false`

Type: Boolean

Required: No

AuthTokenLastModifiedDate

The date the auth token was last modified

Type: Timestamp

Required: No

AutoMinorVersionUpgrade

If you are running Valkey or Redis OSS engine version 6.0 or later, set this parameter to yes if you want to opt-in to the next auto minor version upgrade campaign. This parameter is disabled for previous versions.

Type: Boolean

Required: No

CacheClusterCreateTime

The date and time when the cluster was created.

Type: Timestamp

Required: No

CacheClusterId

The user-supplied identifier of the cluster. This identifier is a unique key that identifies a cluster.

Type: String

Required: No

CacheClusterStatus

The current state of this cluster, one of the following values: `available`, `creating`, `deleted`, `deleting`, `incompatible-network`, `modifying`, `rebooting cluster nodes`, `restore-failed`, or `snapshotting`.

Type: String

Required: No

CacheNodes.CacheNode.N

A list of cache nodes that are members of the cluster.

Type: Array of CacheNode objects

Required: No

CacheNodeType

The name of the compute and memory capacity node type for the cluster.

The following node types are supported by ElastiCache. Generally speaking, the current generation types provide more memory and computational power at lower cost when compared to their equivalent previous generation counterparts.

- General purpose:
  - Current generation:

    M7g node types: `cache.m7g.large`, `cache.m7g.xlarge`, `cache.m7g.2xlarge`, `cache.m7g.4xlarge`, `cache.m7g.8xlarge`, `cache.m7g.12xlarge`, `cache.m7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    M6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.m6g.large`, `cache.m6g.xlarge`, `cache.m6g.2xlarge`, `cache.m6g.4xlarge`, `cache.m6g.8xlarge`, `cache.m6g.12xlarge`, `cache.m6g.16xlarge`

    M5 node types: `cache.m5.large`, `cache.m5.xlarge`, `cache.m5.2xlarge`, `cache.m5.4xlarge`, `cache.m5.12xlarge`, `cache.m5.24xlarge`

    M4 node types: `cache.m4.large`, `cache.m4.xlarge`, `cache.m4.2xlarge`, `cache.m4.4xlarge`, `cache.m4.10xlarge`

    T4g node types (available only for Redis OSS engine version 5.0.6 onward and Memcached engine version 1.5.16 onward): `cache.t4g.micro`, `cache.t4g.small`, `cache.t4g.medium`

    T3 node types: `cache.t3.micro`, `cache.t3.small`, `cache.t3.medium`

T2 node types: `cache.t2.micro`, `cache.t2.small`, `cache.t2.medium`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  T1 node types: `cache.t1.micro`

  M1 node types: `cache.m1.small`, `cache.m1.medium`, `cache.m1.large`, `cache.m1.xlarge`

  M3 node types: `cache.m3.medium`, `cache.m3.large`, `cache.m3.xlarge`, `cache.m3.2xlarge`

- Compute optimized:
  - Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

    C1 node types: `cache.c1.xlarge`

- Memory optimized:
  - Current generation:

    R7g node types: `cache.r7g.large`, `cache.r7g.xlarge`, `cache.r7g.2xlarge`, `cache.r7g.4xlarge`, `cache.r7g.8xlarge`, `cache.r7g.12xlarge`, `cache.r7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    R6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.r6g.large`, `cache.r6g.xlarge`, `cache.r6g.2xlarge`, `cache.r6g.4xlarge`, `cache.r6g.8xlarge`, `cache.r6g.12xlarge`, `cache.r6g.16xlarge`

    R5 node types: `cache.r5.large`, `cache.r5.xlarge`, `cache.r5.2xlarge`, `cache.r5.4xlarge`, `cache.r5.12xlarge`, `cache.r5.24xlarge`

R4 node types: `cache.r4.large`, `cache.r4.xlarge`, `cache.r4.2xlarge`, `cache.r4.4xlarge`, `cache.r4.8xlarge`, `cache.r4.16xlarge`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  M2 node types: `cache.m2.xlarge`, `cache.m2.2xlarge`, `cache.m2.4xlarge`

  R3 node types: `cache.r3.large`, `cache.r3.xlarge`, `cache.r3.2xlarge`, `cache.r3.4xlarge`, `cache.r3.8xlarge`

Additional node type info

- All current generation instance types are created in Amazon VPC by default.
- Valkey or Redis OSS append-only files (AOF) are not supported for T1 or T2 instances.
- Valkey or Redis OSS Multi-AZ with automatic failover is not supported on T1 instances.
- The configuration variables `appendonly` and `appendfsync` are not supported on Valkey, or on Redis OSS version 2.8.22 and later.

Type: String

Required: No

CacheParameterGroup

Status of the cache parameter group.

Type: CacheParameterGroupStatus object

Required: No

CacheSecurityGroups.CacheSecurityGroup.N

A list of cache security group elements, composed of name and status sub-elements.

Type: Array of CacheSecurityGroupMembership objects

Required: No

CacheSubnetGroupName

The name of the cache subnet group associated with the cluster.

Type: String

Required: No

ClientDownloadLandingPage

The URL of the web page where you can download the latest ElastiCache client library.

Type: String

Required: No

ConfigurationEndpoint

Represents a Memcached cluster endpoint which can be used by an application to connect to any node in the cluster. The configuration endpoint will always have `.cfg` in it.

Example: `mem-3.9dvc4r.cfg.usw2.cache.amazonaws.com:11211`

Type: Endpoint object

Required: No

Engine

The name of the cache engine (`memcached` or `redis`) to be used for this cluster.

Type: String

Required: No

EngineVersion

The version of the cache engine that is used in this cluster.

Type: String

Required: No

IpDiscovery

The network type associated with the cluster, either `ipv4` | `ipv6`. IPv6 is supported for workloads using Valkey 7.2 and above, Redis OSS engine version 6.2 to 7.1 or Memcached engine version 1.6.6 and above on all instances built on the Nitro system.

Type: String

Valid Values: `ipv4 | ipv6`

Required: No

LogDeliveryConfigurations.LogDeliveryConfiguration.N

Returns the destination, format and type of the logs.

Type: Array of LogDeliveryConfiguration objects

Required: No

NetworkType

Must be either `ipv4` | `ipv6` | `dual_stack`. IPv6 is supported for workloads using Valkey 7.2 and above, Redis OSS engine version 6.2 7.1 or Memcached engine version 1.6.6 and above on all instances built on the Nitro system.

Type: String

Valid Values: `ipv4 | ipv6 | dual_stack`

Required: No

NotificationConfiguration

Describes a notification topic and its status. Notification topics are used for publishing ElastiCache events to subscribers using Amazon Simple Notification Service (SNS).

Type: NotificationConfiguration object

Required: No

NumCacheNodes

The number of cache nodes in the cluster.

For clusters running Valkey or Redis OSS, this value must be 1. For clusters running Memcached, this value must be between 1 and 40.

Type: Integer

Required: No

PendingModifiedValues

A group of settings that are applied to the cluster in the future, or that are currently being applied.

Type: PendingModifiedValues object

Required: No

PreferredAvailabilityZone

The name of the Availability Zone in which the cluster is located or "Multiple" if the cache nodes are located in different Availability Zones.

Type: String

Required: No

PreferredMaintenanceWindow

Specifies the weekly time range during which maintenance on the cluster is performed. It is specified as a range in the format ddd:hh24:mi-ddd:hh24:mi (24H Clock UTC). The minimum maintenance window is a 60 minute period.

Valid values for ddd are:

- `sun`
- `mon`
- `tue`
- `wed`
- `thu`
- `fri`
- `sat`

Example: `sun:23:00-mon:01:30`

Type: String

Required: No

PreferredOutpostArn

The outpost ARN in which the cache cluster is created.

Type: String

Required: No

ReplicationGroupId

The replication group to which this cluster belongs. If this field is empty, the cluster is not associated with any replication group.

Type: String

Required: No

ReplicationGroupLogDeliveryEnabled

A boolean value indicating whether log delivery is enabled for the replication group.

Type: Boolean

Required: No

SecurityGroups.member.N

A list of VPC Security Groups associated with the cluster.

Type: Array of SecurityGroupMembership objects

Required: No

SnapshotRetentionLimit

The number of days for which ElastiCache retains automatic cluster snapshots before deleting them. For example, if you set `SnapshotRetentionLimit` to 5, a snapshot that was taken today is retained for 5 days before being deleted.

> **Important**
>
> If the value of SnapshotRetentionLimit is set to zero (0), backups are turned off.

Type: Integer

Required: No

SnapshotWindow

The daily time range (in UTC) during which ElastiCache begins taking a daily snapshot of your cluster.

Example: `05:00-09:00`

Type: String

Required: No

TransitEncryptionEnabled

A flag that enables in-transit encryption when set to `true`.

Required: Only available when creating a replication group in an Amazon VPC using Redis OSS version `3.2.6`, `4.x` or later.

Default: `false`

Type: Boolean

Required: No

TransitEncryptionMode

A setting that allows you to migrate your clients to use in-transit encryption, with no downtime.

Type: String

Valid Values: `preferred | required`

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

CacheEngineVersion

Provides all of the details about a particular cache engine version.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheEngineDescription

The description of the cache engine.

Type: String

Required: No

CacheEngineVersionDescription

The description of the cache engine version.

Type: String

Required: No

CacheParameterGroupFamily

The name of the cache parameter group family associated with this cache engine.

Valid values are: `memcached1.4` | `memcached1.5` | `memcached1.6` | `redis2.6` | `redis2.8` | `redis3.2` | `redis4.0` | `redis5.0` | `redis6.x` | `redis7`

Type: String

Required: No

Engine

The name of the cache engine.

Type: String

Required: No

EngineVersion

The version number of the cache engine.

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

CacheNode

Represents an individual cache node within a cluster. Each cache node runs its own instance of the cluster's protocol-compliant caching software - either Memcached, Valkey or Redis OSS.

The following node types are supported by ElastiCache. Generally speaking, the current generation types provide more memory and computational power at lower cost when compared to their equivalent previous generation counterparts.

- General purpose:
  - Current generation:

    M7g node types: `cache.m7g.large`, `cache.m7g.xlarge`, `cache.m7g.2xlarge`, `cache.m7g.4xlarge`, `cache.m7g.8xlarge`, `cache.m7g.12xlarge`, `cache.m7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    M6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.m6g.large`, `cache.m6g.xlarge`, `cache.m6g.2xlarge`, `cache.m6g.4xlarge`, `cache.m6g.8xlarge`, `cache.m6g.12xlarge`, `cache.m6g.16xlarge`

    M5 node types: `cache.m5.large`, `cache.m5.xlarge`, `cache.m5.2xlarge`, `cache.m5.4xlarge`, `cache.m5.12xlarge`, `cache.m5.24xlarge`

    M4 node types: `cache.m4.large`, `cache.m4.xlarge`, `cache.m4.2xlarge`, `cache.m4.4xlarge`, `cache.m4.10xlarge`

    T4g node types (available only for Redis OSS engine version 5.0.6 onward and Memcached engine version 1.5.16 onward): `cache.t4g.micro`, `cache.t4g.small`, `cache.t4g.medium`

    T3 node types: `cache.t3.micro`, `cache.t3.small`, `cache.t3.medium`

    T2 node types: `cache.t2.micro`, `cache.t2.small`, `cache.t2.medium`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  T1 node types: `cache.t1.micro`

  M1 node types: `cache.m1.small`, `cache.m1.medium`, `cache.m1.large`, `cache.m1.xlarge`

  M3 node types: `cache.m3.medium`, `cache.m3.large`, `cache.m3.xlarge`, `cache.m3.2xlarge`

- Compute optimized:
  - Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

    C1 node types: `cache.c1.xlarge`

- Memory optimized:
  - Current generation:

    R7g node types: `cache.r7g.large`, `cache.r7g.xlarge`, `cache.r7g.2xlarge`, `cache.r7g.4xlarge`, `cache.r7g.8xlarge`, `cache.r7g.12xlarge`, `cache.r7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    R6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.r6g.large`, `cache.r6g.xlarge`, `cache.r6g.2xlarge`, `cache.r6g.4xlarge`, `cache.r6g.8xlarge`, `cache.r6g.12xlarge`, `cache.r6g.16xlarge`

    R5 node types: `cache.r5.large`, `cache.r5.xlarge`, `cache.r5.2xlarge`, `cache.r5.4xlarge`, `cache.r5.12xlarge`, `cache.r5.24xlarge`

    R4 node types: `cache.r4.large`, `cache.r4.xlarge`, `cache.r4.2xlarge`, `cache.r4.4xlarge`, `cache.r4.8xlarge`, `cache.r4.16xlarge`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  M2 node types: `cache.m2.xlarge`, `cache.m2.2xlarge`, `cache.m2.4xlarge`

  R3 node types: `cache.r3.large`, `cache.r3.xlarge`, `cache.r3.2xlarge`, `cache.r3.4xlarge`, `cache.r3.8xlarge`

Additional node type info

- All current generation instance types are created in Amazon VPC by default.
- Valkey or Redis OSS append-only files (AOF) are not supported for T1 or T2 instances.
- Valkey or Redis OSS Multi-AZ with automatic failover is not supported on T1 instances.
- The configuration variables `appendonly` and `appendfsync` are not supported on Valkey, or on Redis OSS version 2.8.22 and later.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheNodeCreateTime

The date and time when the cache node was created.

Type: Timestamp

Required: No

CacheNodeId

The cache node identifier. A node ID is a numeric identifier (0001, 0002, etc.). The combination of cluster ID and node ID uniquely identifies every cache node used in a customer's Amazon account.

Type: String

Required: No

CacheNodeStatus

The current state of this cache node, one of the following values: `available`, `creating`, `rebooting`, or `deleting`.

Type: String

Required: No

CustomerAvailabilityZone

The Availability Zone where this node was created and now resides.

Type: String

Required: No

CustomerOutpostArn

The customer outpost ARN of the cache node.

Type: String

Required: No

Endpoint

The hostname for connecting to this cache node.

Type: Endpoint object

Required: No

ParameterGroupStatus

The status of the parameter group applied to this cache node.

Type: String

Required: No

SourceCacheNodeId

The ID of the primary node to which this read replica node is synchronized. If this field is empty, this node is not associated with a primary cluster.

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

CacheNodeTypeSpecificParameter

A parameter that has a different value for each cache node type it is applied to. For example, in a Valkey or Redis OSS cluster, a `cache.m1.large` cache node type would have a larger `maxmemory` value than a `cache.m1.small` type.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

AllowedValues

The valid range of values for the parameter.

Type: String

Required: No

CacheNodeTypeSpecificValues.CacheNodeTypeSpecificValue.N

A list of cache node types and their corresponding values for this parameter.

Type: Array of CacheNodeTypeSpecificValue objects

Required: No

ChangeType

Indicates whether a change to the parameter is applied immediately or requires a reboot for the change to be applied. You can force a reboot or wait until the next maintenance window's reboot. For more information, see Rebooting a Cluster.

Type: String

Valid Values: `immediate | requires-reboot`

Required: No

DataType

The valid data type for the parameter.

Type: String

Required: No

Description

A description of the parameter.

Type: String

Required: No

IsModifiable

Indicates whether (`true`) or not (`false`) the parameter can be modified. Some parameters have security or operational implications that prevent them from being changed.

Type: Boolean

Required: No

MinimumEngineVersion

The earliest cache engine version to which the parameter can apply.

Type: String

Required: No

ParameterName

The name of the parameter.

Type: String

Required: No

Source

The source of the parameter value.

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

CacheNodeTypeSpecificValue

A value that applies only to a certain cache node type.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheNodeType

The cache node type for which this value applies.

Type: String

Required: No

Value

The value for the cache node 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

CacheNodeUpdateStatus

The status of the service update on the cache node

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheNodeId

The node ID of the cache cluster

Type: String

Required: No

NodeDeletionDate

The deletion date of the node

Type: Timestamp

Required: No

NodeUpdateEndDate

The end date of the update for a node

Type: Timestamp

Required: No

NodeUpdateInitiatedBy

Reflects whether the update was initiated by the customer or automatically applied

Type: String

Valid Values: `system | customer`

Required: No

NodeUpdateInitiatedDate

The date when the update is triggered

Type: Timestamp

Required: No

NodeUpdateStartDate

The start date of the update for a node

Type: Timestamp

Required: No

NodeUpdateStatus

The update status of the node

Type: String

Valid Values: `not-applied | waiting-to-start | in-progress | stopping | stopped | complete`

Required: No

NodeUpdateStatusModifiedDate

The date when the NodeUpdateStatus was last modified>

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

CacheParameterGroup

Represents the output of a `CreateCacheParameterGroup` operation.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

ARN

The ARN (Amazon Resource Name) of the cache parameter group.

Type: String

Required: No

CacheParameterGroupFamily

The name of the cache parameter group family that this cache parameter group is compatible with.

Valid values are: `memcached1.4` | `memcached1.5` | `memcached1.6` | `redis2.6` | `redis2.8` | `redis3.2` | `redis4.0` | `redis5.0` | `redis6.x` | `redis7`

Type: String

Required: No

CacheParameterGroupName

The name of the cache parameter group.

Type: String

Required: No

Description

The description for this cache parameter group.

Type: String

Required: No

IsGlobal

Indicates whether the parameter group is associated with a Global datastore

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

CacheParameterGroupStatus

Status of the cache parameter group.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheNodeIdsToReboot.CacheNodeId.N

A list of the cache node IDs which need to be rebooted for parameter changes to be applied. A node ID is a numeric identifier (0001, 0002, etc.).

Type: Array of strings

Required: No

CacheParameterGroupName

The name of the cache parameter group.

Type: String

Required: No

ParameterApplyStatus

The status of parameter updates.

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

CacheSecurityGroup

Represents the output of one of the following operations:

- `AuthorizeCacheSecurityGroupIngress`
- `CreateCacheSecurityGroup`
- `RevokeCacheSecurityGroupIngress`

Contents

> **Note**
>
> In the following list, the required parameters are described first.

ARN

The ARN of the cache security group,

Type: String

Required: No

CacheSecurityGroupName

The name of the cache security group.

Type: String

Required: No

Description

The description of the cache security group.

Type: String

Required: No

EC2SecurityGroups.EC2SecurityGroup.N

A list of Amazon EC2 security groups that are associated with this cache security group.

Type: Array of EC2SecurityGroup objects

Required: No

OwnerId

The Amazon account ID of the cache security group owner.

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

CacheSecurityGroupMembership

Represents a cluster's status within a particular cache security group.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheSecurityGroupName

The name of the cache security group.

Type: String

Required: No

Status

The membership status in the cache security group. The status changes when a cache security group is modified, or when the cache security groups assigned to a cluster are modified.

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

CacheSubnetGroup

Represents the output of one of the following operations:

- `CreateCacheSubnetGroup`
- `ModifyCacheSubnetGroup`

Contents

> **Note**
> In the following list, the required parameters are described first.

ARN

The ARN (Amazon Resource Name) of the cache subnet group.

Type: String

Required: No

CacheSubnetGroupDescription

The description of the cache subnet group.

Type: String

Required: No

CacheSubnetGroupName

The name of the cache subnet group.

Type: String

Required: No

Subnets.Subnet.N

A list of subnets associated with the cache subnet group.

Type: Array of Subnet objects

Required: No

SupportedNetworkTypes.member.N

Either `ipv4` | `ipv6` | `dual_stack`. IPv6 is supported for workloads using Valkey 7.2 and above, Redis OSS engine version 6.2 to 7.1 or Memcached engine version 1.6.6 and above on all instances built on the Nitro system.

Type: Array of strings

Valid Values: `ipv4 | ipv6 | dual_stack`

Required: No

VpcId

The Amazon Virtual Private Cloud identifier (VPC ID) of the cache subnet group.

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

CacheUsageLimits

The usage limits for storage and ElastiCache Processing Units for the cache.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

DataStorage

The maximum data storage limit in the cache, expressed in Gigabytes.

Type: DataStorage object

Required: No

ECPUPerSecond

The configuration for the number of ElastiCache Processing Units (ECPU) the cache can consume per second.

Type: ECPUPerSecond 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

CloudWatchLogsDestinationDetails

The configuration details of the CloudWatch Logs destination.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

LogGroup

The name of the CloudWatch Logs log group.

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

ConfigureShard

Node group (shard) configuration options when adding or removing replicas. Each node group (shard) configuration has the following members: NodeGroupId, NewReplicaCount, and PreferredAvailabilityZones.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

NewReplicaCount

The number of replicas you want in this node group at the end of this operation. The maximum value for `NewReplicaCount` is 5. The minimum value depends upon the type of Valkey or Redis OSS replication group you are working with.

The minimum number of replicas in a shard or replication group is:

- Valkey or Redis OSS (cluster mode disabled)
  - If Multi-AZ: 1
  - If Multi-AZ: 0
- Valkey or Redis OSS (cluster mode enabled): 0 (though you will not be able to failover to a replica if your primary node fails)

Type: Integer

Required: Yes

NodeGroupId

The 4-digit id for the node group you are configuring. For Valkey or Redis OSS (cluster mode disabled) replication groups, the node group id is always 0001. To find a Valkey or Redis OSS (cluster mode enabled)'s node group's (shard's) id, see Finding a Shard's Id.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 4.

Pattern: \d+

Required: Yes

PreferredAvailabilityZones.PreferredAvailabilityZone.N

A list of `PreferredAvailabilityZone` strings that specify which availability zones the replication group's nodes are to be in. The nummber of `PreferredAvailabilityZone` values must equal the value of `NewReplicaCount` plus 1 to account for the primary node. If this member of `ReplicaConfiguration` is omitted, ElastiCache selects the availability zone for each of the replicas.

Type: Array of strings

Required: No

PreferredOutpostArns.PreferredOutpostArn.N

The outpost ARNs in which the cache cluster is created.

Type: Array of strings

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

CustomerNodeEndpoint

The endpoint from which data should be migrated.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

Address

The address of the node endpoint

Type: String

Required: No

Port

The port of the node endpoint

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

DataStorage

The data storage limit.

Contents

> **Note**
> In the following list, the required parameters are described first.

Unit

The unit that the storage is measured in, in GB.

Type: String

Valid Values: GB

Required: Yes

Maximum

The upper limit for data storage the cache is set to use.

Type: Integer

Required: No

Minimum

The lower limit for data storage the cache is set to use.

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

DestinationDetails

Configuration details of either a CloudWatch Logs destination or Kinesis Data Firehose destination.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CloudWatchLogsDetails

The configuration details of the CloudWatch Logs destination.

Type: CloudWatchLogsDestinationDetails object

Required: No

KinesisFirehoseDetails

The configuration details of the Kinesis Data Firehose destination.

Type: KinesisFirehoseDestinationDetails 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

EC2SecurityGroup

Provides ownership and status information for an Amazon EC2 security group.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

EC2SecurityGroupName

The name of the Amazon EC2 security group.

Type: String

Required: No

EC2SecurityGroupOwnerId

The Amazon account ID of the Amazon EC2 security group owner.

Type: String

Required: No

Status

The status of the Amazon EC2 security group.

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

ECPUPerSecond

The configuration for the number of ElastiCache Processing Units (ECPU) the cache can consume per second.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

Maximum

The configuration for the maximum number of ECPUs the cache can consume per second.

Type: Integer

Required: No

Minimum

The configuration for the minimum number of ECPUs the cache should be able consume per second.

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

Endpoint

Represents the information required for client programs to connect to a cache node. This value is read-only.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

Address

The DNS hostname of the cache node.

Type: String

Required: No

Port

The port number that the cache engine is listening on.

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

EngineDefaults

Represents the output of a `DescribeEngineDefaultParameters` operation.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheNodeTypeSpecificParameters.CacheNodeTypeSpecificParameter.N

A list of parameters specific to a particular cache node type. Each element in the list contains detailed information about one parameter.

Type: Array of CacheNodeTypeSpecificParameter objects

Required: No

CacheParameterGroupFamily

Specifies the name of the cache parameter group family to which the engine default parameters apply.

Valid values are: `memcached1.4` | `memcached1.5` | `memcached1.6` | `redis2.6` | `redis2.8` | `redis3.2` | `redis4.0` | `redis5.0` | `redis6.0` | `redis6.x` | `redis7`

Type: String

Required: No

Marker

Provides an identifier to allow retrieval of paginated results.

Type: String

Required: No

Parameters.Parameter.N

Contains a list of engine default parameters.

Type: Array of Parameter 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

Event

Represents a single occurrence of something interesting within the system. Some examples of events are creating a cluster, adding or removing a cache node, or rebooting a node.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

Date

The date and time when the event occurred.

Type: Timestamp

Required: No

Message

The text of the event.

Type: String

Required: No

SourceIdentifier

The identifier for the source of the event. For example, if the event occurred at the cluster level, the identifier would be the name of the cluster.

Type: String

Required: No

SourceType

Specifies the origin of this event - a cluster, a parameter group, a security group, etc.

Type: String

Valid Values: `cache-cluster | cache-parameter-group | cache-security-group | cache-subnet-group | replication-group | serverless-cache | serverless-cache-snapshot | user | user-group`

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

Filter

Used to streamline results of a search based on the property being filtered.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

Name

The property being filtered. For example, UserId.

Type: String

Pattern: `.*\S.*`

Required: Yes

Values.member.N

The property values to filter on. For example, "user-123".

Type: Array of strings

Array Members: Minimum number of 1 item.

Pattern: `.*\S.*`

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

GlobalNodeGroup

Indicates the slot configuration and global identifier for a slice group.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

GlobalNodeGroupId

The name of the global node group

Type: String

Required: No

Slots

The keyspace for this node group

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

GlobalReplicationGroup

Consists of a primary cluster that accepts writes and an associated secondary cluster that resides in a different Amazon region. The secondary cluster accepts only reads. The primary cluster automatically replicates updates to the secondary cluster.

- The **GlobalReplicationGroupIdSuffix** represents the name of the Global datastore, which is what you use to associate a secondary cluster.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

ARN

The ARN (Amazon Resource Name) of the global replication group.

Type: String

Required: No

AtRestEncryptionEnabled

A flag that enables encryption at rest when set to `true`.

You cannot modify the value of `AtRestEncryptionEnabled` after the replication group is created. To enable encryption at rest on a replication group you must set `AtRestEncryptionEnabled` to `true` when you create the replication group.

Required: Only available when creating a replication group in an Amazon VPC using Redis OSS version `3.2.6`, `4.x` or later.

Type: Boolean

Required: No

AuthTokenEnabled

A flag that enables using an `AuthToken` (password) when issuing Valkey or Redis OSS commands.

Default: `false`

Type: Boolean

Required: No

CacheNodeType

The cache node type of the Global datastore

Type: String

Required: No

ClusterEnabled

A flag that indicates whether the Global datastore is cluster enabled.

Type: Boolean

Required: No

Engine

The ElastiCache engine. For Valkey or Redis OSS only.

Type: String

Required: No

EngineVersion

The ElastiCache engine version.

Type: String

Required: No

GlobalNodeGroups.GlobalNodeGroup.N

Indicates the slot configuration and global identifier for each slice group.

Type: Array of GlobalNodeGroup objects

Required: No

GlobalReplicationGroupDescription

The optional description of the Global datastore

Type: String

Required: No

GlobalReplicationGroupId

The name of the Global datastore

Type: String

Required: No

Members.GlobalReplicationGroupMember.N

The replication groups that comprise the Global datastore.

Type: Array of GlobalReplicationGroupMember objects

Required: No

Status

The status of the Global datastore

Type: String

Required: No

TransitEncryptionEnabled

A flag that enables in-transit encryption when set to true.

Required: Only available when creating a replication group in an Amazon VPC using Redis OSS version `3.2.6`, `4.x` or later.

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

GlobalReplicationGroupInfo

The name of the Global datastore and role of this replication group in the Global datastore.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

GlobalReplicationGroupId

The name of the Global datastore

Type: String

Required: No

GlobalReplicationGroupMemberRole

The role of the replication group in a Global datastore. Can be primary or secondary.

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

GlobalReplicationGroupMember

A member of a Global datastore. It contains the Replication Group Id, the Amazon region and the role of the replication group.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

AutomaticFailover

Indicates whether automatic failover is enabled for the replication group.

Type: String

Valid Values: `enabled | disabled | enabling | disabling`

Required: No

ReplicationGroupId

The replication group id of the Global datastore member.

Type: String

Required: No

ReplicationGroupRegion

The Amazon region of the Global datastore member.

Type: String

Required: No

Role

Indicates the role of the replication group, primary or secondary.

Type: String

Required: No

Status

The status of the membership of the replication group.

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

KinesisFirehoseDestinationDetails

The configuration details of the Kinesis Data Firehose destination.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

DeliveryStream

The name of the Kinesis Data Firehose delivery stream.

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

LogDeliveryConfiguration

Returns the destination, format and type of the logs.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

DestinationDetails

Configuration details of either a CloudWatch Logs destination or Kinesis Data Firehose destination.

Type: DestinationDetails object

Required: No

DestinationType

Returns the destination type, either `cloudwatch-logs` or `kinesis-firehose`.

Type: String

Valid Values: `cloudwatch-logs | kinesis-firehose`

Required: No

LogFormat

Returns the log format, either JSON or TEXT.

Type: String

Valid Values: `text | json`

Required: No

LogType

Refers to slow-log or engine-log.

Type: String

Valid Values: `slow-log | engine-log`

Required: No

Message

Returns an error message for the log delivery configuration.

Type: String

Required: No

Status

Returns the log delivery configuration status. Values are one of `enabling` | `disabling` | `modifying` | `active` | `error`

Type: String

Valid Values: `active | enabling | modifying | disabling | error`

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

LogDeliveryConfigurationRequest

Specifies the destination, format and type of the logs.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

DestinationDetails

Configuration details of either a CloudWatch Logs destination or Kinesis Data Firehose destination.

Type: DestinationDetails object

Required: No

DestinationType

Specify either `cloudwatch-logs` or `kinesis-firehose` as the destination type.

Type: String

Valid Values: `cloudwatch-logs | kinesis-firehose`

Required: No

Enabled

Specify if log delivery is enabled. Default `true`.

Type: Boolean

Required: No

LogFormat

Specifies either JSON or TEXT

Type: String

Valid Values: `text | json`

Required: No

LogType

Refers to slow-log or engine-log..

Type: String

Valid Values: `slow-log | engine-log`

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

NodeGroup

Represents a collection of cache nodes in a replication group. One node in the node group is the read/write primary node. All the other nodes are read-only Replica nodes.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

NodeGroupId

The identifier for the node group (shard). A Valkey or Redis OSS (cluster mode disabled) replication group contains only 1 node group; therefore, the node group ID is 0001. A Valkey or Redis OSS (cluster mode enabled) replication group contains 1 to 90 node groups numbered 0001 to 0090. Optionally, the user can provide the id for a node group.

Type: String

Required: No

NodeGroupMembers.NodeGroupMember.N

A list containing information about individual nodes within the node group (shard).

Type: Array of NodeGroupMember objects

Required: No

PrimaryEndpoint

The endpoint of the primary node in this node group (shard).

Type: Endpoint object

Required: No

ReaderEndpoint

The endpoint of the replica nodes in this node group (shard). This value is read-only.

Type: Endpoint object

Required: No

Slots

The keyspace for this node group (shard).

Type: String

Required: No

Status

The current state of this replication group - `creating`, `available`, `modifying`, `deleting`.

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

NodeGroupConfiguration

Node group (shard) configuration options. Each node group (shard) configuration has the following: `Slots`, `PrimaryAvailabilityZone`, `ReplicaAvailabilityZones`, `ReplicaCount`.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

NodeGroupId

Either the ElastiCache supplied 4-digit id or a user supplied id for the node group these configuration values apply to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 4.

Pattern: \d+

Required: No

PrimaryAvailabilityZone

The Availability Zone where the primary node of this node group (shard) is launched.

Type: String

Required: No

PrimaryOutpostArn

The outpost ARN of the primary node.

Type: String

Required: No

ReplicaAvailabilityZones.AvailabilityZone.N

A list of Availability Zones to be used for the read replicas. The number of Availability Zones in this list must match the value of `ReplicaCount` or `ReplicasPerNodeGroup` if not specified.

Type: Array of strings

Required: No

ReplicaCount

The number of read replica nodes in this node group (shard).

Type: Integer

Required: No

ReplicaOutpostArns.OutpostArn.N

The outpost ARN of the node replicas.

Type: Array of strings

Required: No

Slots

A string that specifies the keyspace for a particular node group. Keyspaces range from 0 to 16,383. The string is in the format `startkey-endkey`.

Example: `"0-3999"`

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

NodeGroupMember

Represents a single node within a node group (shard).

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheClusterId

The ID of the cluster to which the node belongs.

Type: String

Required: No

CacheNodeId

The ID of the node within its cluster. A node ID is a numeric identifier (0001, 0002, etc.).

Type: String

Required: No

CurrentRole

The role that is currently assigned to the node - `primary` or `replica`. This member is only applicable for Valkey or Redis OSS (cluster mode disabled) replication groups.

Type: String

Required: No

PreferredAvailabilityZone

The name of the Availability Zone in which the node is located.

Type: String

Required: No

PreferredOutpostArn

The outpost ARN of the node group member.

Type: String

Required: No

ReadEndpoint

The information required for client programs to connect to a node for read operations. The read endpoint is only applicable on Valkey or Redis OSS (cluster mode disabled) clusters.

Type: Endpoint 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

NodeGroupMemberUpdateStatus

The status of the service update on the node group member

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheClusterId

The cache cluster ID

Type: String

Required: No

CacheNodeId

The node ID of the cache cluster

Type: String

Required: No

NodeDeletionDate

The deletion date of the node

Type: Timestamp

Required: No

NodeUpdateEndDate

The end date of the update for a node

Type: Timestamp

Required: No

NodeUpdateInitiatedBy

Reflects whether the update was initiated by the customer or automatically applied

Type: String

Valid Values: `system | customer`

Required: No

NodeUpdateInitiatedDate

The date when the update is triggered

Type: Timestamp

Required: No

NodeUpdateStartDate

The start date of the update for a node

Type: Timestamp

Required: No

NodeUpdateStatus

The update status of the node

Type: String

Valid Values: `not-applied | waiting-to-start | in-progress | stopping | stopped | complete`

Required: No

NodeUpdateStatusModifiedDate

The date when the NodeUpdateStatus was last modified

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

NodeGroupUpdateStatus

The status of the service update on the node group

Contents

> **Note**
>
> In the following list, the required parameters are described first.

NodeGroupId

The ID of the node group

Type: String

Required: No

NodeGroupMemberUpdateStatus.NodeGroupMemberUpdateStatus.N

The status of the service update on the node group member

Type: Array of NodeGroupMemberUpdateStatus 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

NodeSnapshot

Represents an individual cache node in a snapshot of a cluster.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheClusterId

A unique identifier for the source cluster.

Type: String

Required: No

CacheNodeCreateTime

The date and time when the cache node was created in the source cluster.

Type: Timestamp

Required: No

CacheNodeId

The cache node identifier for the node in the source cluster.

Type: String

Required: No

CacheSize

The size of the cache on the source cache node.

Type: String

Required: No

NodeGroupConfiguration

The configuration for the source node group (shard).

Type: NodeGroupConfiguration object

Required: No

NodeGroupId

A unique identifier for the source node group (shard).

Type: String

Required: No

SnapshotCreateTime

The date and time when the source node's metadata and cache data set was obtained for the snapshot.

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

NotificationConfiguration

Describes a notification topic and its status. Notification topics are used for publishing ElastiCache events to subscribers using Amazon Simple Notification Service (SNS).

Contents

> **Note**
>
> In the following list, the required parameters are described first.

TopicArn

The Amazon Resource Name (ARN) that identifies the topic.

Type: String

Required: No

TopicStatus

The current state of the topic.

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

Parameter

Describes an individual setting that controls some aspect of ElastiCache behavior.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

AllowedValues

The valid range of values for the parameter.

Type: String

Required: No

ChangeType

Indicates whether a change to the parameter is applied immediately or requires a reboot for the change to be applied. You can force a reboot or wait until the next maintenance window's reboot. For more information, see Rebooting a Cluster.

Type: String

Valid Values: `immediate | requires-reboot`

Required: No

DataType

The valid data type for the parameter.

Type: String

Required: No

Description

A description of the parameter.

Type: String

Required: No

IsModifiable

Indicates whether (`true`) or not (`false`) the parameter can be modified. Some parameters have security or operational implications that prevent them from being changed.

Type: Boolean

Required: No

MinimumEngineVersion

The earliest cache engine version to which the parameter can apply.

Type: String

Required: No

ParameterName

The name of the parameter.

Type: String

Required: No

ParameterValue

The value of the parameter.

Type: String

Required: No

Source

The source of the parameter.

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

ParameterNameValue

Describes a name-value pair that is used to update the value of a parameter.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

ParameterName

The name of the parameter.

Type: String

Required: No

ParameterValue

The value of the parameter.

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

PendingLogDeliveryConfiguration

The log delivery configurations being modified

Contents

> **Note**
>
> In the following list, the required parameters are described first.

DestinationDetails

Configuration details of either a CloudWatch Logs destination or Kinesis Data Firehose destination.

Type: DestinationDetails object

Required: No

DestinationType

Returns the destination type, either CloudWatch Logs or Kinesis Data Firehose.

Type: String

Valid Values: `cloudwatch-logs | kinesis-firehose`

Required: No

LogFormat

Returns the log format, either JSON or TEXT

Type: String

Valid Values: `text | json`

Required: No

LogType

Refers to slow-log or engine-log..

Type: String

Valid Values: `slow-log | engine-log`

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

PendingModifiedValues

A group of settings that are applied to the cluster in the future, or that are currently being applied.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

AuthTokenStatus

The auth token status

Type: String

Valid Values: `SETTING | ROTATING`

Required: No

CacheNodeIdsToRemove.CacheNodeId.N

A list of cache node IDs that are being removed (or will be removed) from the cluster. A node ID is a 4-digit numeric identifier (0001, 0002, etc.).

Type: Array of strings

Required: No

CacheNodeType

The cache node type that this cluster or replication group is scaled to.

Type: String

Required: No

EngineVersion

The new cache engine version that the cluster runs.

Type: String

Required: No

NumCacheNodes

The new number of cache nodes for the cluster.

For clusters running Valkey or Redis OSS, this value must be 1. For clusters running Memcached, this value must be between 1 and 40.

Type: Integer

Required: No

PendingLogDeliveryConfiguration.member.N

The log delivery configurations being modified

Type: Array of PendingLogDeliveryConfiguration objects

Required: No

ScaleConfig

The scaling configuration changes that are pending for the Memcached cluster.

Type: ScaleConfig object

Required: No

TransitEncryptionEnabled

A flag that enables in-transit encryption when set to true.

Type: Boolean

Required: No

TransitEncryptionMode

A setting that allows you to migrate your clients to use in-transit encryption, with no downtime.

Type: String

Valid Values: `preferred | required`

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

ProcessedUpdateAction

Update action that has been processed for the corresponding apply/stop request

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheClusterId

The ID of the cache cluster

Type: String

Required: No

ReplicationGroupId

The ID of the replication group

Type: String

Required: No

ServiceUpdateName

The unique ID of the service update

Type: String

Required: No

UpdateActionStatus

The status of the update action on the Valkey or Redis OSS cluster

Type: String

Valid Values: `not-applied | waiting-to-start | in-progress | stopping | stopped | complete | scheduling | scheduled | not-applicable`

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

RecurringCharge

Contains the specific price and frequency of a recurring charges for a reserved cache node, or for a reserved cache node offering.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

RecurringChargeAmount

The monetary amount of the recurring charge.

Type: Double

Required: No

RecurringChargeFrequency

The frequency of the recurring charge.

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

RegionalConfiguration

A list of the replication groups

Contents

> **Note**
> In the following list, the required parameters are described first.

ReplicationGroupId

The name of the secondary cluster

Type: String

Required: Yes

ReplicationGroupRegion

The Amazon region where the cluster is stored

Type: String

Required: Yes

ReshardingConfiguration.ReshardingConfiguration.N

A list of `PreferredAvailabilityZones` objects that specifies the configuration of a node group in the resharded cluster.

Type: Array of ReshardingConfiguration objects

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

ReplicationGroup

Contains all of the attributes of a specific Valkey or Redis OSS replication group.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

ARN

The ARN (Amazon Resource Name) of the replication group.

Type: String

Required: No

AtRestEncryptionEnabled

A flag that enables encryption at-rest when set to `true`.

You cannot modify the value of `AtRestEncryptionEnabled` after the cluster is created. To enable encryption at-rest on a cluster you must set `AtRestEncryptionEnabled` to `true` when you create a cluster.

Required: Only available when creating a replication group in an Amazon VPC using Redis OSS version `3.2.6`, `4.x` or later.

Default: `false`

Type: Boolean

Required: No

AuthTokenEnabled

A flag that enables using an `AuthToken` (password) when issuing Valkey or Redis OSS commands.

Default: `false`

Type: Boolean

Required: No

AuthTokenLastModifiedDate

The date the auth token was last modified

Type: Timestamp

Required: No

AutomaticFailover

Indicates the status of automatic failover for this Valkey or Redis OSS replication group.

Type: String

Valid Values: `enabled | disabled | enabling | disabling`

Required: No

AutoMinorVersionUpgrade

If you are running Valkey 7.2 and above, or Redis OSS engine version 6.0 and above, set this parameter to yes if you want to opt-in to the next auto minor version upgrade campaign. This parameter is disabled for previous versions.

Type: Boolean

Required: No

CacheNodeType

The name of the compute and memory capacity node type for each node in the replication group.

Type: String

Required: No

ClusterEnabled

A flag indicating whether or not this replication group is cluster enabled; i.e., whether its data can be partitioned across multiple shards (API/CLI: node groups).

Valid values: `true` | `false`

Type: Boolean

Required: No

ClusterMode

Enabled or Disabled. To modify cluster mode from Disabled to Enabled, you must first set the cluster mode to Compatible. Compatible mode allows your Valkey or Redis OSS clients to connect using both cluster mode enabled and cluster mode disabled. After you migrate all Valkey or Redis OSS clients to use cluster mode enabled, you can then complete cluster mode configuration and set the cluster mode to Enabled.

Type: String

Valid Values: `enabled | disabled | compatible`

Required: No

ConfigurationEndpoint

The configuration endpoint for this replication group. Use the configuration endpoint to connect to this replication group.

Type: Endpoint object

Required: No

DataTiering

Enables data tiering. Data tiering is only supported for replication groups using the r6gd node type. This parameter must be set to true when using r6gd nodes. For more information, see Data tiering.

Type: String

Valid Values: `enabled | disabled`

Required: No

Description

The user supplied description of the replication group.

Type: String

Required: No

Engine

The engine used in a replication group. The options are valkey, memcached or redis.

Type: String

Required: No

GlobalReplicationGroupInfo

The name of the Global datastore and role of this replication group in the Global datastore.

Type: GlobalReplicationGroupInfo object

Required: No

IpDiscovery

The network type you choose when modifying a cluster, either `ipv4 | ipv6`. IPv6 is supported for workloads using Valkey 7.2 and above, Redis OSS engine version 6.2 to 7.1 or Memcached engine version 1.6.6 and above on all instances built on the Nitro system.

Type: String

Valid Values: `ipv4 | ipv6`

Required: No

KmsKeyId

The ID of the KMS key used to encrypt the disk in the cluster.

Type: String

Required: No

LogDeliveryConfigurations.LogDeliveryConfiguration.N

Returns the destination, format and type of the logs.

Type: Array of LogDeliveryConfiguration objects

Required: No

MemberClusters.ClusterId.N

The names of all the cache clusters that are part of this replication group.

Type: Array of strings

Required: No

MemberClustersOutpostArns.ReplicationGroupOutpostArn.N

The outpost ARNs of the replication group's member clusters.

Type: Array of strings

Required: No

MultiAZ

A flag indicating if you have Multi-AZ enabled to enhance fault tolerance. For more information, see Minimizing Downtime: Multi-AZ

Type: String

Valid Values: `enabled | disabled`

Required: No

NetworkType

Must be either `ipv4` | `ipv6` | `dual_stack`. IPv6 is supported for workloads using Valkey 7.2 and above, Redis OSS engine version 6.2 to 7.1 or Memcached engine version 1.6.6 and above on all instances built on the Nitro system.

Type: String

Valid Values: `ipv4 | ipv6 | dual_stack`

Required: No

NodeGroups.NodeGroup.N

A list of node groups in this replication group. For Valkey or Redis OSS (cluster mode disabled) replication groups, this is a single-element list. For Valkey or Redis OSS (cluster mode enabled) replication groups, the list contains an entry for each node group (shard).

Type: Array of NodeGroup objects

Required: No

PendingModifiedValues

A group of settings to be applied to the replication group, either immediately or during the next maintenance window.

Type: ReplicationGroupPendingModifiedValues object

Required: No

ReplicationGroupCreateTime

The date and time when the cluster was created.

Type: Timestamp

Required: No

ReplicationGroupId

The identifier for the replication group.

Type: String

Required: No

SnapshotRetentionLimit

The number of days for which ElastiCache retains automatic cluster snapshots before deleting them. For example, if you set `SnapshotRetentionLimit` to 5, a snapshot that was taken today is retained for 5 days before being deleted.

> **Important**
>
> If the value of `SnapshotRetentionLimit` is set to zero (0), backups are turned off.

Type: Integer

Required: No

SnapshottingClusterId

The cluster ID that is used as the daily snapshot source for the replication group.

Type: String

Required: No

SnapshotWindow

The daily time range (in UTC) during which ElastiCache begins taking a daily snapshot of your node group (shard).

Example: `05:00-09:00`

If you do not specify this parameter, ElastiCache automatically chooses an appropriate time range.

> **Note**
>
> This parameter is only valid if the `Engine` parameter is `redis`.

Type: String

Required: No

Status

The current state of this replication group - `creating`, `available`, `modifying`, `deleting`, `create-failed`, `snapshotting`.

Type: String

Required: No

TransitEncryptionEnabled

A flag that enables in-transit encryption when set to `true`.

Required: Only available when creating a replication group in an Amazon VPC using Redis OSS version `3.2.6`, `4.x` or later.

Default: `false`

Type: Boolean

Required: No

TransitEncryptionMode

A setting that allows you to migrate your clients to use in-transit encryption, with no downtime.

Type: String

Valid Values: `preferred | required`

Required: No

UserGroupIds.member.N

The ID of the user group associated to the replication group.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

ReplicationGroupPendingModifiedValues

The settings to be applied to the Valkey or Redis OSS replication group, either immediately or during the next maintenance window.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

AuthTokenStatus

The auth token status

Type: String

Valid Values: `SETTING | ROTATING`

Required: No

AutomaticFailoverStatus

Indicates the status of automatic failover for this Valkey or Redis OSS replication group.

Type: String

Valid Values: `enabled | disabled`

Required: No

ClusterMode

Enabled or Disabled. To modify cluster mode from Disabled to Enabled, you must first set the cluster mode to Compatible. Compatible mode allows your Valkey or Redis OSS clients to connect using both cluster mode enabled and cluster mode disabled. After you migrate all Valkey or Redis OSS clients to use cluster mode enabled, you can then complete cluster mode configuration and set the cluster mode to Enabled.

Type: String

Valid Values: `enabled | disabled | compatible`

Required: No

PendingLogDeliveryConfiguration.member.N

The log delivery configurations being modified

Type: Array of PendingLogDeliveryConfiguration objects

Required: No

PrimaryClusterId

The primary cluster ID that is applied immediately (if `--apply-immediately` was specified), or during the next maintenance window.

Type: String

Required: No

Resharding

The status of an online resharding operation.

Type: ReshardingStatus object

Required: No

TransitEncryptionEnabled

A flag that enables in-transit encryption when set to true.

Type: Boolean

Required: No

TransitEncryptionMode

A setting that allows you to migrate your clients to use in-transit encryption, with no downtime.

Type: String

Valid Values: `preferred | required`

Required: No

UserGroups

The user group being modified.

Type: UserGroupsUpdateStatus 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

ReservedCacheNode

Represents the output of a `PurchaseReservedCacheNodesOffering` operation.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheNodeCount

The number of cache nodes that have been reserved.

Type: Integer

Required: No

CacheNodeType

The cache node type for the reserved cache nodes.

The following node types are supported by ElastiCache. Generally speaking, the current generation types provide more memory and computational power at lower cost when compared to their equivalent previous generation counterparts.

- General purpose:
  - Current generation:

    M7g node types: `cache.m7g.large`, `cache.m7g.xlarge`, `cache.m7g.2xlarge`, `cache.m7g.4xlarge`, `cache.m7g.8xlarge`, `cache.m7g.12xlarge`, `cache.m7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    M6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.m6g.large`, `cache.m6g.xlarge`,

`cache.m6g.2xlarge`, `cache.m6g.4xlarge`, `cache.m6g.8xlarge`, `cache.m6g.12xlarge`, `cache.m6g.16xlarge`

M5 node types: `cache.m5.large`, `cache.m5.xlarge`, `cache.m5.2xlarge`, `cache.m5.4xlarge`, `cache.m5.12xlarge`, `cache.m5.24xlarge`

M4 node types: `cache.m4.large`, `cache.m4.xlarge`, `cache.m4.2xlarge`, `cache.m4.4xlarge`, `cache.m4.10xlarge`

T4g node types (available only for Redis OSS engine version 5.0.6 onward and Memcached engine version 1.5.16 onward): `cache.t4g.micro`, `cache.t4g.small`, `cache.t4g.medium`

T3 node types: `cache.t3.micro`, `cache.t3.small`, `cache.t3.medium`

T2 node types: `cache.t2.micro`, `cache.t2.small`, `cache.t2.medium`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  T1 node types: `cache.t1.micro`

  M1 node types: `cache.m1.small`, `cache.m1.medium`, `cache.m1.large`, `cache.m1.xlarge`

  M3 node types: `cache.m3.medium`, `cache.m3.large`, `cache.m3.xlarge`, `cache.m3.2xlarge`

- Compute optimized:
  - Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

    C1 node types: `cache.c1.xlarge`

- Memory optimized:
  - Current generation:

    R7g node types: `cache.r7g.large`, `cache.r7g.xlarge`, `cache.r7g.2xlarge`, `cache.r7g.4xlarge`, `cache.r7g.8xlarge`, `cache.r7g.12xlarge`, `cache.r7g.16xlarge`

> **Note**
>
> For region availability, see Supported Node Types

R6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.r6g.large`, `cache.r6g.xlarge`, `cache.r6g.2xlarge`, `cache.r6g.4xlarge`, `cache.r6g.8xlarge`, `cache.r6g.12xlarge`, `cache.r6g.16xlarge`

R5 node types: `cache.r5.large`, `cache.r5.xlarge`, `cache.r5.2xlarge`, `cache.r5.4xlarge`, `cache.r5.12xlarge`, `cache.r5.24xlarge`

R4 node types: `cache.r4.large`, `cache.r4.xlarge`, `cache.r4.2xlarge`, `cache.r4.4xlarge`, `cache.r4.8xlarge`, `cache.r4.16xlarge`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  M2 node types: `cache.m2.xlarge`, `cache.m2.2xlarge`, `cache.m2.4xlarge`

  R3 node types: `cache.r3.large`, `cache.r3.xlarge`, `cache.r3.2xlarge`, `cache.r3.4xlarge`, `cache.r3.8xlarge`

Additional node type info

- All current generation instance types are created in Amazon VPC by default.
- Valkey or Redis OSS append-only files (AOF) are not supported for T1 or T2 instances.
- Valkey or Redis OSS Multi-AZ with automatic failover is not supported on T1 instances.
- The configuration variables `appendonly` and `appendfsync` are not supported on Valkey, or on Redis OSS version 2.8.22 and later.

Type: String

Required: No

Duration

The duration of the reservation in seconds.

Type: Integer

Required: No

FixedPrice

The fixed price charged for this reserved cache node.

Type: Double

Required: No

OfferingType

The offering type of this reserved cache node.

Type: String

Required: No

ProductDescription

The description of the reserved cache node.

Type: String

Required: No

RecurringCharges.RecurringCharge.N

The recurring price charged to run this reserved cache node.

Type: Array of RecurringCharge objects

Required: No

ReservationARN

The Amazon Resource Name (ARN) of the reserved cache node.

Example: `arn:aws:elasticache:us-east-1:123456789012:reserved-instance:ri-2017-03-27-08-33-25-582`

Type: String

Required: No

ReservedCacheNodeId

The unique identifier for the reservation.

Type: String

Required: No

ReservedCacheNodesOfferingId

The offering identifier.

Type: String

Required: No

StartTime

The time the reservation started.

Type: Timestamp

Required: No

State

The state of the reserved cache node.

Type: String

Required: No

UsagePrice

The hourly price charged for this reserved cache node.

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

ReservedCacheNodesOffering

Describes all of the attributes of a reserved cache node offering.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheNodeType

The cache node type for the reserved cache node.

The following node types are supported by ElastiCache. Generally speaking, the current generation types provide more memory and computational power at lower cost when compared to their equivalent previous generation counterparts.

- General purpose:
  - Current generation:

    M7g node types: `cache.m7g.large`, `cache.m7g.xlarge`, `cache.m7g.2xlarge`, `cache.m7g.4xlarge`, `cache.m7g.8xlarge`, `cache.m7g.12xlarge`, `cache.m7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    M6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.m6g.large`, `cache.m6g.xlarge`, `cache.m6g.2xlarge`, `cache.m6g.4xlarge`, `cache.m6g.8xlarge`, `cache.m6g.12xlarge`, `cache.m6g.16xlarge`

    M5 node types: `cache.m5.large`, `cache.m5.xlarge`, `cache.m5.2xlarge`, `cache.m5.4xlarge`, `cache.m5.12xlarge`, `cache.m5.24xlarge`

M4 node types: `cache.m4.large`, `cache.m4.xlarge`, `cache.m4.2xlarge`, `cache.m4.4xlarge`, `cache.m4.10xlarge`

T4g node types (available only for Redis OSS engine version 5.0.6 onward and Memcached engine version 1.5.16 onward): `cache.t4g.micro`, `cache.t4g.small`, `cache.t4g.medium`

T3 node types: `cache.t3.micro`, `cache.t3.small`, `cache.t3.medium`

T2 node types: `cache.t2.micro`, `cache.t2.small`, `cache.t2.medium`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  T1 node types: `cache.t1.micro`

  M1 node types: `cache.m1.small`, `cache.m1.medium`, `cache.m1.large`, `cache.m1.xlarge`

  M3 node types: `cache.m3.medium`, `cache.m3.large`, `cache.m3.xlarge`, `cache.m3.2xlarge`

- Compute optimized:
  - Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

    C1 node types: `cache.c1.xlarge`

- Memory optimized:
  - Current generation:

    R7g node types: `cache.r7g.large`, `cache.r7g.xlarge`, `cache.r7g.2xlarge`, `cache.r7g.4xlarge`, `cache.r7g.8xlarge`, `cache.r7g.12xlarge`, `cache.r7g.16xlarge`

> **Note**
>
> For region availability, see Supported Node Types

R6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.r6g.large`, `cache.r6g.xlarge`, `cache.r6g.2xlarge`, `cache.r6g.4xlarge`, `cache.r6g.8xlarge`, `cache.r6g.12xlarge`, `cache.r6g.16xlarge`

R5 node types: `cache.r5.large`, `cache.r5.xlarge`, `cache.r5.2xlarge`, `cache.r5.4xlarge`, `cache.r5.12xlarge`, `cache.r5.24xlarge`

R4 node types: `cache.r4.large`, `cache.r4.xlarge`, `cache.r4.2xlarge`, `cache.r4.4xlarge`, `cache.r4.8xlarge`, `cache.r4.16xlarge`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  M2 node types: `cache.m2.xlarge`, `cache.m2.2xlarge`, `cache.m2.4xlarge`

  R3 node types: `cache.r3.large`, `cache.r3.xlarge`, `cache.r3.2xlarge`, `cache.r3.4xlarge`, `cache.r3.8xlarge`

Additional node type info

- All current generation instance types are created in Amazon VPC by default.
- Valkey or Redis OSS append-only files (AOF) are not supported for T1 or T2 instances.
- Valkey or Redis OSS Multi-AZ with automatic failover is not supported on T1 instances.
- The configuration variables `appendonly` and `appendfsync` are not supported on Valkey, or on Redis OSS version 2.8.22 and later.

Type: String

Required: No

Duration

The duration of the offering. in seconds.

Type: Integer

Required: No

FixedPrice

The fixed price charged for this offering.

Type: Double

Required: No

OfferingType

The offering type.

Type: String

Required: No

ProductDescription

The cache engine used by the offering.

Type: String

Required: No

RecurringCharges.RecurringCharge.N

The recurring price charged to run this reserved cache node.

Type: Array of RecurringCharge objects

Required: No

ReservedCacheNodesOfferingId

A unique identifier for the reserved cache node offering.

Type: String

Required: No

UsagePrice

The hourly price charged for this offering.

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

ReshardingConfiguration

A list of `PreferredAvailabilityZones` objects that specifies the configuration of a node group in the resharded cluster.

Contents

> **Note**
> In the following list, the required parameters are described first.

NodeGroupId

Either the ElastiCache supplied 4-digit id or a user supplied id for the node group these configuration values apply to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 4.

Pattern: \d+

Required: No

PreferredAvailabilityZones.AvailabilityZone.N

A list of preferred availability zones for the nodes in this cluster.

Type: Array of strings

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

ReshardingStatus

The status of an online resharding operation.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

SlotMigration

Represents the progress of an online resharding operation.

Type: SlotMigration 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

ScaleConfig

Configuration settings for horizontal or vertical scaling operations on Memcached clusters.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

ScaleIntervalMinutes

The time interval in seconds between scaling operations when performing gradual scaling for a Memcached cluster.

Type: Integer

Required: No

ScalePercentage

The percentage by which to scale the Memcached cluster, either horizontally by adding nodes or vertically by increasing resources.

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

SecurityGroupMembership

Represents a single cache security group and its status.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

SecurityGroupId

The identifier of the cache security group.

Type: String

Required: No

Status

The status of the cache security group membership. The status changes whenever a cache security group is modified, or when the cache security groups assigned to a cluster are modified.

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

ServerlessCache

The resource representing a serverless cache.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

ARN

The Amazon Resource Name (ARN) of the serverless cache.

Type: String

Required: No

CacheUsageLimits

The cache usage limit for the serverless cache.

Type: CacheUsageLimits object

Required: No

CreateTime

When the serverless cache was created.

Type: Timestamp

Required: No

DailySnapshotTime

The daily time that a cache snapshot will be created. Default is NULL, i.e. snapshots will not be created at a specific time on a daily basis. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: No

Description

A description of the serverless cache.

Type: String

Required: No

Endpoint

Represents the information required for client programs to connect to a cache node. This value is read-only.

Type: Endpoint object

Required: No

Engine

The engine the serverless cache is compatible with.

Type: String

Required: No

FullEngineVersion

The name and version number of the engine the serverless cache is compatible with.

Type: String

Required: No

KmsKeyId

The ID of the AWS Key Management Service (KMS) key that is used to encrypt data at rest in the serverless cache.

Type: String

Required: No

MajorEngineVersion

The version number of the engine the serverless cache is compatible with.

Type: String

Required: No

ReaderEndpoint

Represents the information required for client programs to connect to a cache node. This value is read-only.

Type: Endpoint object

Required: No

SecurityGroupIds.SecurityGroupId.N

The IDs of the EC2 security groups associated with the serverless cache.

Type: Array of strings

Required: No

ServerlessCacheName

The unique identifier of the serverless cache.

Type: String

Required: No

SnapshotRetentionLimit

The current setting for the number of serverless cache snapshots the system will retain. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: Integer

Required: No

Status

The current status of the serverless cache. The allowed values are CREATING, AVAILABLE, DELETING, CREATE-FAILED and MODIFYING.

Type: String

Required: No

SubnetIds.SubnetId.N

If no subnet IDs are given and your VPC is in us-west-1, then ElastiCache will select 2 default subnets across AZs in your VPC. For all other Regions, if no subnet IDs are given then ElastiCache will select 3 default subnets across AZs in your default VPC.

Type: Array of strings

Required: No

UserGroupId

The identifier of the user group associated with the serverless cache. Available for Valkey and Redis OSS only. Default is NULL.

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

ServerlessCacheConfiguration

The configuration settings for a specific serverless cache.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

Engine

The engine that the serverless cache is configured with.

Type: String

Required: No

MajorEngineVersion

The engine version number that the serverless cache is configured with.

Type: String

Required: No

ServerlessCacheName

The identifier of a serverless cache.

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

ServerlessCacheSnapshot

The resource representing a serverless cache snapshot. Available for Valkey, Redis OSS and Serverless Memcached only.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

ARN

The Amazon Resource Name (ARN) of a serverless cache snapshot. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: No

BytesUsedForCache

The total size of a serverless cache snapshot, in bytes. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: No

CreateTime

The date and time that the source serverless cache's metadata and cache data set was obtained for the snapshot. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: Timestamp

Required: No

ExpiryTime

The time that the serverless cache snapshot will expire. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: Timestamp

Required: No

KmsKeyId

The ID of the AWS Key Management Service (KMS) key of a serverless cache snapshot. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: No

ServerlessCacheConfiguration

The configuration of the serverless cache, at the time the snapshot was taken. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: ServerlessCacheConfiguration object

Required: No

ServerlessCacheSnapshotName

The identifier of a serverless cache snapshot. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: No

SnapshotType

The type of snapshot of serverless cache. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: String

Required: No

Status

The current status of the serverless cache. Available for Valkey, Redis OSS and Serverless Memcached only.

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

ServiceUpdate

An update that you can apply to your Valkey or Redis OSS clusters.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

AutoUpdateAfterRecommendedApplyByDate

Indicates whether the service update will be automatically applied once the recommended apply-by date has expired.

Type: Boolean

Required: No

Engine

The Elasticache engine to which the update applies. Either Valkey, Redis OSS or Memcached.

Type: String

Required: No

EngineVersion

The Elasticache engine version to which the update applies. Either Valkey, Redis OSS or Memcached engine version.

Type: String

Required: No

EstimatedUpdateTime

The estimated length of time the service update will take

Type: String

Required: No

ServiceUpdateDescription

Provides details of the service update

Type: String

Required: No

ServiceUpdateEndDate

The date after which the service update is no longer available

Type: Timestamp

Required: No

ServiceUpdateName

The unique ID of the service update

Type: String

Required: No

ServiceUpdateRecommendedApplyByDate

The recommendend date to apply the service update in order to ensure compliance. For information on compliance, see Self-Service Security Updates for Compliance.

Type: Timestamp

Required: No

ServiceUpdateReleaseDate

The date when the service update is initially available

Type: Timestamp

Required: No

ServiceUpdateSeverity

The severity of the service update

Type: String

Valid Values: `critical | important | medium | low`

Required: No

ServiceUpdateStatus

The status of the service update

Type: String

Valid Values: `available | cancelled | expired`

Required: No

ServiceUpdateType

Reflects the nature of the service update

Type: String

Valid Values: `security-update`

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

SlotMigration

Represents the progress of an online resharding operation.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

ProgressPercentage

The percentage of the slot migration that is complete.

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

Snapshot

Represents a copy of an entire Valkey or Redis OSS cluster as of the time when the snapshot was taken.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

ARN

The ARN (Amazon Resource Name) of the snapshot.

Type: String

Required: No

AutomaticFailover

Indicates the status of automatic failover for the source Valkey or Redis OSS replication group.

Type: String

Valid Values: `enabled | disabled | enabling | disabling`

Required: No

AutoMinorVersionUpgrade

If you are running Valkey 7.2 and above or Redis OSS engine version 6.0 and above, set this parameter to yes if you want to opt-in to the next auto minor version upgrade campaign. This parameter is disabled for previous versions.

Type: Boolean

Required: No

CacheClusterCreateTime

The date and time when the source cluster was created.

Type: Timestamp

Required: No

CacheClusterId

The user-supplied identifier of the source cluster.

Type: String

Required: No

CacheNodeType

The name of the compute and memory capacity node type for the source cluster.

The following node types are supported by ElastiCache. Generally speaking, the current generation types provide more memory and computational power at lower cost when compared to their equivalent previous generation counterparts.

- General purpose:
  - Current generation:

    M7g node types: `cache.m7g.large`, `cache.m7g.xlarge`, `cache.m7g.2xlarge`, `cache.m7g.4xlarge`, `cache.m7g.8xlarge`, `cache.m7g.12xlarge`, `cache.m7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    M6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.m6g.large`, `cache.m6g.xlarge`, `cache.m6g.2xlarge`, `cache.m6g.4xlarge`, `cache.m6g.8xlarge`, `cache.m6g.12xlarge`, `cache.m6g.16xlarge`

    M5 node types: `cache.m5.large`, `cache.m5.xlarge`, `cache.m5.2xlarge`, `cache.m5.4xlarge`, `cache.m5.12xlarge`, `cache.m5.24xlarge`

    M4 node types: `cache.m4.large`, `cache.m4.xlarge`, `cache.m4.2xlarge`, `cache.m4.4xlarge`, `cache.m4.10xlarge`

T4g node types (available only for Redis OSS engine version 5.0.6 onward and Memcached engine version 1.5.16 onward): `cache.t4g.micro`, `cache.t4g.small`, `cache.t4g.medium`

T3 node types: `cache.t3.micro`, `cache.t3.small`, `cache.t3.medium`

T2 node types: `cache.t2.micro`, `cache.t2.small`, `cache.t2.medium`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  T1 node types: `cache.t1.micro`

  M1 node types: `cache.m1.small`, `cache.m1.medium`, `cache.m1.large`, `cache.m1.xlarge`

  M3 node types: `cache.m3.medium`, `cache.m3.large`, `cache.m3.xlarge`, `cache.m3.2xlarge`

- Compute optimized:
  - Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

    C1 node types: `cache.c1.xlarge`

- Memory optimized:
  - Current generation:

    R7g node types: `cache.r7g.large`, `cache.r7g.xlarge`, `cache.r7g.2xlarge`, `cache.r7g.4xlarge`, `cache.r7g.8xlarge`, `cache.r7g.12xlarge`, `cache.r7g.16xlarge`

    > **Note**
    >
    > For region availability, see Supported Node Types

    R6g node types (available only for Redis OSS engine version 5.0.6 onward and for Memcached engine version 1.5.16 onward): `cache.r6g.large`, `cache.r6g.xlarge`,

`cache.r6g.2xlarge`, `cache.r6g.4xlarge`, `cache.r6g.8xlarge`, `cache.r6g.12xlarge`, `cache.r6g.16xlarge`

R5 node types: `cache.r5.large`, `cache.r5.xlarge`, `cache.r5.2xlarge`, `cache.r5.4xlarge`, `cache.r5.12xlarge`, `cache.r5.24xlarge`

R4 node types: `cache.r4.large`, `cache.r4.xlarge`, `cache.r4.2xlarge`, `cache.r4.4xlarge`, `cache.r4.8xlarge`, `cache.r4.16xlarge`

- Previous generation: (not recommended. Existing clusters are still supported but creation of new clusters is not supported for these types.)

  M2 node types: `cache.m2.xlarge`, `cache.m2.2xlarge`, `cache.m2.4xlarge`

  R3 node types: `cache.r3.large`, `cache.r3.xlarge`, `cache.r3.2xlarge`, `cache.r3.4xlarge`, `cache.r3.8xlarge`

Additional node type info

- All current generation instance types are created in Amazon VPC by default.
- Valkey or Redis OSS append-only files (AOF) are not supported for T1 or T2 instances.
- Valkey or Redis OSS Multi-AZ with automatic failover is not supported on T1 instances.
- The configuration variables `appendonly` and `appendfsync` are not supported on Valkey, or on Redis OSS version 2.8.22 and later.

Type: String

Required: No

CacheParameterGroupName

The cache parameter group that is associated with the source cluster.

Type: String

Required: No

CacheSubnetGroupName

The name of the cache subnet group associated with the source cluster.

Type: String

Required: No

DataTiering

Enables data tiering. Data tiering is only supported for replication groups using the r6gd node type. This parameter must be set to true when using r6gd nodes. For more information, see Data tiering.

Type: String

Valid Values: `enabled | disabled`

Required: No

Engine

The name of the cache engine (`memcached` or `redis`) used by the source cluster.

Type: String

Required: No

EngineVersion

The version of the cache engine version that is used by the source cluster.

Type: String

Required: No

KmsKeyId

The ID of the KMS key used to encrypt the snapshot.

Type: String

Required: No

NodeSnapshots.NodeSnapshot.N

A list of the cache nodes in the source cluster.

Type: Array of NodeSnapshot objects

Required: No

NumCacheNodes

The number of cache nodes in the source cluster.

For clusters running Valkey or Redis OSS, this value must be 1. For clusters running Memcached, this value must be between 1 and 40.

Type: Integer

Required: No

NumNodeGroups

The number of node groups (shards) in this snapshot. When restoring from a snapshot, the number of node groups (shards) in the snapshot and in the restored replication group must be the same.

Type: Integer

Required: No

Port

The port number used by each cache nodes in the source cluster.

Type: Integer

Required: No

PreferredAvailabilityZone

The name of the Availability Zone in which the source cluster is located.

Type: String

Required: No

PreferredMaintenanceWindow

Specifies the weekly time range during which maintenance on the cluster is performed. It is specified as a range in the format ddd:hh24:mi-ddd:hh24:mi (24H Clock UTC). The minimum maintenance window is a 60 minute period.

Valid values for ddd are:

- `sun`
- `mon`
- `tue`

- `wed`
- `thu`
- `fri`
- `sat`

Example: `sun:23:00-mon:01:30`

Type: String

Required: No

PreferredOutpostArn

The ARN (Amazon Resource Name) of the preferred outpost.

Type: String

Required: No

ReplicationGroupDescription

A description of the source replication group.

Type: String

Required: No

ReplicationGroupId

The unique identifier of the source replication group.

Type: String

Required: No

SnapshotName

The name of a snapshot. For an automatic snapshot, the name is system-generated. For a manual snapshot, this is the user-provided name.

Type: String

Required: No

SnapshotRetentionLimit

For an automatic snapshot, the number of days for which ElastiCache retains the snapshot before deleting it.

For manual snapshots, this field reflects the `SnapshotRetentionLimit` for the source cluster when the snapshot was created. This field is otherwise ignored: Manual snapshots do not expire, and can only be deleted using the `DeleteSnapshot` operation.

Important If the value of SnapshotRetentionLimit is set to zero (0), backups are turned off.

Type: Integer

Required: No

SnapshotSource

Indicates whether the snapshot is from an automatic backup (`automated`) or was created manually (`manual`).

Type: String

Required: No

SnapshotStatus

The status of the snapshot. Valid values: `creating` | `available` | `restoring` | `copying` | `deleting`.

Type: String

Required: No

SnapshotWindow

The daily time range during which ElastiCache takes daily snapshots of the source cluster.

Type: String

Required: No

TopicArn

The Amazon Resource Name (ARN) for the topic used by the source cluster for publishing notifications.

Type: String

Required: No

VpcId

The Amazon Virtual Private Cloud identifier (VPC ID) of the cache subnet group for the source cluster.

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

Subnet

Represents the subnet associated with a cluster. This parameter refers to subnets defined in Amazon Virtual Private Cloud (Amazon VPC) and used with ElastiCache.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

SubnetAvailabilityZone

The Availability Zone associated with the subnet.

Type: AvailabilityZone object

Required: No

SubnetIdentifier

The unique identifier for the subnet.

Type: String

Required: No

SubnetOutpost

The outpost ARN of the subnet.

Type: SubnetOutpost object

Required: No

SupportedNetworkTypes.member.N

Either `ipv4` | `ipv6` | `dual_stack`. IPv6 is supported for workloads using Valkey 7.2 and above, Redis OSS engine version 6.2 to 7.1 or Memcached engine version 1.6.6 and above on all instances built on the Nitro system.

Type: Array of strings

Valid Values: `ipv4 | ipv6 | dual_stack`

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

SubnetOutpost

The ID of the outpost subnet.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

SubnetOutpostArn

The outpost ARN of the subnet.

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

Tag

A tag that can be added to an ElastiCache cluster or replication group. Tags are composed of a Key/Value pair. You can use tags to categorize and track all your ElastiCache resources, with the exception of global replication group. When you add or remove tags on replication groups, those actions will be replicated to all nodes in the replication group. A tag with a null Value is permitted.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

Key

The key for the tag. May not be null.

Type: String

Required: No

Value

The tag's value. May be null.

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

TimeRangeFilter

Filters update actions from the service updates that are in available status during the time range.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

EndTime

The end time of the time range filter

Type: Timestamp

Required: No

StartTime

The start time of the time range filter

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

UnprocessedUpdateAction

Update action that has failed to be processed for the corresponding apply/stop request

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheClusterId

The ID of the cache cluster

Type: String

Required: No

ErrorMessage

The error message that describes the reason the request was not processed

Type: String

Required: No

ErrorType

The error type for requests that are not processed

Type: String

Required: No

ReplicationGroupId

The replication group ID

Type: String

Required: No

ServiceUpdateName

The unique ID of the service update

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

UpdateAction

The status of the service update for a specific replication group

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CacheClusterId

The ID of the cache cluster

Type: String

Required: No

CacheNodeUpdateStatus.CacheNodeUpdateStatus.N

The status of the service update on the cache node

Type: Array of CacheNodeUpdateStatus objects

Required: No

Engine

The Elasticache engine to which the update applies. Either Valkey, Redis OSS or Memcached.

Type: String

Required: No

EstimatedUpdateTime

The estimated length of time for the update to complete

Type: String

Required: No

NodeGroupUpdateStatus.NodeGroupUpdateStatus.N

The status of the service update on the node group

Type: Array of NodeGroupUpdateStatus objects

Required: No

NodesUpdated

The progress of the service update on the replication group

Type: String

Required: No

ReplicationGroupId

The ID of the replication group

Type: String

Required: No

ServiceUpdateName

The unique ID of the service update

Type: String

Required: No

ServiceUpdateRecommendedApplyByDate

The recommended date to apply the service update to ensure compliance. For information on compliance, see Self-Service Security Updates for Compliance.

Type: Timestamp

Required: No

ServiceUpdateReleaseDate

The date the update is first available

Type: Timestamp

Required: No

ServiceUpdateSeverity

The severity of the service update

Type: String

Valid Values: `critical | important | medium | low`

Required: No

ServiceUpdateStatus

The status of the service update

Type: String

Valid Values: `available | cancelled | expired`

Required: No

ServiceUpdateType

Reflects the nature of the service update

Type: String

Valid Values: `security-update`

Required: No

SlaMet

If yes, all nodes in the replication group have been updated by the recommended apply-by date. If no, at least one node in the replication group have not been updated by the recommended apply-by date. If N/A, the replication group was created after the recommended apply-by date.

Type: String

Valid Values: `yes | no | n/a`

Required: No

UpdateActionAvailableDate

The date that the service update is available to a replication group

Type: Timestamp

Required: No

UpdateActionStatus

The status of the update action

Type: String

Valid Values: `not-applied | waiting-to-start | in-progress | stopping | stopped | complete | scheduling | scheduled | not-applicable`

Required: No

UpdateActionStatusModifiedDate

The date when the UpdateActionStatus was last modified

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

User

Contents

> **Note**
>
> In the following list, the required parameters are described first.

AccessString

Access permissions string used for this user.

Type: String

Required: No

ARN

The Amazon Resource Name (ARN) of the user.

Type: String

Required: No

Authentication

Denotes whether the user requires a password to authenticate.

Type: Authentication object

Required: No

Engine

The options are valkey or redis.

Type: String

Pattern: `[a-zA-Z]*`

Required: No

MinimumEngineVersion

The minimum engine version required, which is Redis OSS 6.0

Type: String

Required: No

Status

Indicates the user status. Can be "active", "modifying" or "deleting".

Type: String

Required: No

UserGroupIds.member.N

Returns a list of the user group IDs the user belongs to.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

UserId

The ID of the user.

Type: String

Required: No

UserName

The username of the user.

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

UserGroup

Contents

> **Note**
>
> In the following list, the required parameters are described first.

ARN

The Amazon Resource Name (ARN) of the user group.

Type: String

Required: No

Engine

The options are valkey or redis.

Type: String

Pattern: `[a-zA-Z]*`

Required: No

MinimumEngineVersion

The minimum engine version required, which is Redis OSS 6.0

Type: String

Required: No

PendingChanges

A list of updates being applied to the user group.

Type: UserGroupPendingChanges object

Required: No

ReplicationGroups.member.N

A list of replication groups that the user group can access.

Type: Array of strings

Required: No

ServerlessCaches.member.N

Indicates which serverless caches the specified user group is associated with. Available for Valkey, Redis OSS and Serverless Memcached only.

Type: Array of strings

Required: No

Status

Indicates user group status. Can be "creating", "active", "modifying", "deleting".

Type: String

Required: No

UserGroupId

The ID of the user group.

Type: String

Required: No

UserIds.member.N

The list of user IDs that belong to the user group.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

UserGroupPendingChanges

Returns the updates being applied to the user group.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

UserIdsToAdd.member.N

The list of user IDs to add.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

UserIdsToRemove.member.N

The list of user IDs to remove.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

See Also

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

- AWS SDK for C++

- AWS SDK for Java V2
- AWS SDK for Ruby V3

UserGroupsUpdateStatus

The status of the user group update.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

UserGroupIdsToAdd.member.N

The ID of the user group to add.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

UserGroupIdsToRemove.member.N

The ID of the user group to remove.

Type: Array of strings

Length Constraints: Minimum length of 1.

Pattern: `[a-zA-Z][a-zA-Z0-9\-]*`

Required: No

See Also

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

- AWS SDK for C++

- AWS SDK for Java V2
- AWS SDK for Ruby V3

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: 400

IncompleteSignature

The request signature does not conform to AWS standards.

HTTP Status Code: 400

InternalFailure

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

HTTP Status Code: 500

InvalidAction

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

HTTP Status Code: 400

InvalidClientTokenId

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

HTTP Status Code: 403

InvalidParameterCombination

Parameters that must not be used together were used together.

HTTP Status Code: 400

InvalidParameterValue

An invalid or out-of-range value was supplied for the input parameter.

HTTP Status Code: 400

InvalidQueryParameter

The AWS query string is malformed or does not adhere to AWS standards.

HTTP Status Code: 400

MalformedQueryString

The query string contains a syntax error.

HTTP Status Code: 404

MissingAction

The request is missing an action or a required parameter.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 403

MissingParameter

A required parameter for the specified action is not supplied.

HTTP Status Code: 400

NotAuthorized

You do not have permission to perform this action.

HTTP Status Code: 400

OptInRequired

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

HTTP Status Code: 403

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

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

ValidationError

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

HTTP Status Code: 400