



API Reference

# Amazon Neptune



**API Version 2014-10-31**

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

# Amazon Neptune: API Reference

Copyright © 2024 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

<b>Welcome</b> .....	<b>1</b>
<b>Actions</b> .....	<b>2</b>
AddRoleToDBCluster .....	5
Request Parameters .....	5
Errors .....	5
See Also .....	6
AddSourceIdentifierToSubscription .....	7
Request Parameters .....	7
Response Elements .....	7
Errors .....	8
See Also .....	8
AddTagsToResource .....	9
Request Parameters .....	9
Errors .....	9
See Also .....	10
ApplyPendingMaintenanceAction .....	11
Request Parameters .....	11
Response Elements .....	12
Errors .....	12
See Also .....	12
CopyDBClusterParameterGroup .....	13
Request Parameters .....	13
Response Elements .....	14
Errors .....	14
See Also .....	15
CopyDBClusterSnapshot .....	16
Request Parameters .....	16
Response Elements .....	18
Errors .....	18
See Also .....	19
CopyDBParameterGroup .....	20
Request Parameters .....	20
Response Elements .....	21
Errors .....	21

---

See Also .....	22
CreateDBCluster .....	23
Request Parameters .....	23
Response Elements .....	30
Errors .....	30
See Also .....	33
CreateDBClusterEndpoint .....	34
Request Parameters .....	34
Response Elements .....	35
Errors .....	36
See Also .....	37
CreateDBClusterParameterGroup .....	39
Request Parameters .....	39
Response Elements .....	40
Errors .....	41
See Also .....	41
CreateDBClusterSnapshot .....	42
Request Parameters .....	42
Response Elements .....	43
Errors .....	43
See Also .....	44
CreateDBInstance .....	45
Request Parameters .....	45
Response Elements .....	55
Errors .....	55
See Also .....	58
CreateDBParameterGroup .....	59
Request Parameters .....	59
Response Elements .....	60
Errors .....	61
See Also .....	61
CreateDBSubnetGroup .....	62
Request Parameters .....	62
Response Elements .....	63
Errors .....	63
See Also .....	64

CreateEventSubscription .....	65
Request Parameters .....	65
Response Elements .....	67
Errors .....	67
See Also .....	68
CreateGlobalCluster .....	69
Request Parameters .....	69
Response Elements .....	70
Errors .....	71
See Also .....	71
DeleteDBCluster .....	73
Request Parameters .....	73
Response Elements .....	74
Errors .....	74
See Also .....	75
DeleteDBClusterEndpoint .....	76
Request Parameters .....	76
Response Elements .....	76
Errors .....	77
See Also .....	78
DeleteDBClusterParameterGroup .....	79
Request Parameters .....	79
Errors .....	79
See Also .....	80
DeleteDBClusterSnapshot .....	81
Request Parameters .....	81
Response Elements .....	81
Errors .....	81
See Also .....	82
DeleteDBInstance .....	83
Request Parameters .....	83
Response Elements .....	84
Errors .....	85
See Also .....	85
DeleteDBParameterGroup .....	87
Request Parameters .....	87

---

Errors .....	87
See Also .....	88
DeleteDBSubnetGroup .....	89
Request Parameters .....	89
Errors .....	89
See Also .....	90
DeleteEventSubscription .....	91
Request Parameters .....	91
Response Elements .....	91
Errors .....	91
See Also .....	92
DeleteGlobalCluster .....	93
Request Parameters .....	93
Response Elements .....	93
Errors .....	93
See Also .....	94
DescribeDBClusterEndpoints .....	95
Request Parameters .....	95
Response Elements .....	96
Errors .....	97
See Also .....	97
DescribeDBClusterParameterGroups .....	98
Request Parameters .....	98
Response Elements .....	99
Errors .....	99
See Also .....	99
DescribeDBClusterParameters .....	101
Request Parameters .....	101
Response Elements .....	102
Errors .....	102
See Also .....	103
DescribeDBClusters .....	104
Request Parameters .....	104
Response Elements .....	105
Errors .....	106
See Also .....	106

DescribeDBClusterSnapshotAttributes .....	107
Request Parameters .....	107
Response Elements .....	107
Errors .....	108
See Also .....	108
DescribeDBClusterSnapshots .....	109
Request Parameters .....	109
Response Elements .....	111
Errors .....	112
See Also .....	112
DescribeDBEngineVersions .....	113
Request Parameters .....	113
Response Elements .....	115
Errors .....	115
See Also .....	115
DescribeDBInstances .....	117
Request Parameters .....	117
Response Elements .....	118
Errors .....	119
See Also .....	119
DescribeDBParameterGroups .....	120
Request Parameters .....	120
Response Elements .....	121
Errors .....	121
See Also .....	121
DescribeDBParameters .....	123
Request Parameters .....	123
Response Elements .....	124
Errors .....	124
See Also .....	125
DescribeDBSubnetGroups .....	126
Request Parameters .....	126
Response Elements .....	127
Errors .....	127
See Also .....	127
DescribeEngineDefaultClusterParameters .....	129

Request Parameters .....	129
Response Elements .....	130
Errors .....	130
See Also .....	130
<b>DescribeEngineDefaultParameters .....</b>	<b>131</b>
Request Parameters .....	131
Response Elements .....	132
Errors .....	132
See Also .....	132
<b>DescribeEventCategories .....</b>	<b>133</b>
Request Parameters .....	133
Response Elements .....	133
Errors .....	133
See Also .....	134
<b>DescribeEvents .....</b>	<b>135</b>
Request Parameters .....	135
Response Elements .....	137
Errors .....	137
See Also .....	138
<b>DescribeEventSubscriptions .....</b>	<b>139</b>
Request Parameters .....	139
Response Elements .....	140
Errors .....	140
See Also .....	140
<b>DescribeGlobalClusters .....</b>	<b>142</b>
Request Parameters .....	142
Response Elements .....	143
Errors .....	143
See Also .....	143
<b>DescribeOrderableDBInstanceOptions .....</b>	<b>145</b>
Request Parameters .....	145
Response Elements .....	146
Errors .....	147
See Also .....	147
<b>DescribePendingMaintenanceActions .....</b>	<b>148</b>
Request Parameters .....	148



Response Elements .....	149
Errors .....	149
See Also .....	150
DescribeValidDBInstanceModifications .....	151
Request Parameters .....	151
Response Elements .....	151
Errors .....	151
See Also .....	152
FailoverDBCluster .....	153
Request Parameters .....	153
Response Elements .....	153
Errors .....	154
See Also .....	154
FailoverGlobalCluster .....	156
Request Parameters .....	156
Response Elements .....	157
Errors .....	157
See Also .....	158
ListTagsForResource .....	159
Request Parameters .....	159
Response Elements .....	159
Errors .....	159
See Also .....	160
ModifyDBCluster .....	161
Request Parameters .....	161
Response Elements .....	166
Errors .....	167
See Also .....	168
ModifyDBClusterEndpoint .....	170
Request Parameters .....	170
Response Elements .....	170
Errors .....	172
See Also .....	173
ModifyDBClusterParameterGroup .....	174
Request Parameters .....	174
Response Elements .....	175

Errors .....	175
See Also .....	176
ModifyDBClusterSnapshotAttribute .....	177
Request Parameters .....	177
Response Elements .....	178
Errors .....	179
See Also .....	179
ModifyDBInstance .....	180
Request Parameters .....	180
Response Elements .....	189
Errors .....	190
See Also .....	192
ModifyDBParameterGroup .....	193
Request Parameters .....	193
Response Elements .....	194
Errors .....	194
See Also .....	195
ModifyDBSubnetGroup .....	196
Request Parameters .....	196
Response Elements .....	196
Errors .....	197
See Also .....	198
ModifyEventSubscription .....	199
Request Parameters .....	199
Response Elements .....	200
Errors .....	200
See Also .....	201
ModifyGlobalCluster .....	202
Request Parameters .....	202
Response Elements .....	203
Errors .....	204
See Also .....	204
PromoteReadReplicaDBCluster .....	206
Request Parameters .....	206
Response Elements .....	206
Errors .....	206

See Also .....	207
RebootDBInstance .....	208
Request Parameters .....	208
Response Elements .....	208
Errors .....	209
See Also .....	209
RemoveFromGlobalCluster .....	210
Request Parameters .....	210
Response Elements .....	210
Errors .....	211
See Also .....	211
RemoveRoleFromDBCluster .....	212
Request Parameters .....	212
Errors .....	212
See Also .....	213
RemoveSourceIdentifierFromSubscription .....	214
Request Parameters .....	214
Response Elements .....	214
Errors .....	214
See Also .....	215
RemoveTagsFromResource .....	216
Request Parameters .....	216
Errors .....	216
See Also .....	217
ResetDBClusterParameterGroup .....	218
Request Parameters .....	218
Response Elements .....	219
Errors .....	219
See Also .....	220
ResetDBParameterGroup .....	221
Request Parameters .....	221
Response Elements .....	222
Errors .....	222
See Also .....	222
RestoreDBClusterFromSnapshot .....	224
Request Parameters .....	224

Response Elements .....	228
Errors .....	229
See Also .....	231
RestoreDBClusterToPointInTime .....	233
Request Parameters .....	233
Response Elements .....	238
Errors .....	238
See Also .....	240
StartDBCluster .....	242
Request Parameters .....	242
Response Elements .....	242
Errors .....	242
See Also .....	243
StopDBCluster .....	244
Request Parameters .....	244
Response Elements .....	244
Errors .....	244
See Also .....	245
<b>Data Types .....</b>	<b>246</b>
AvailabilityZone .....	248
Contents .....	248
See Also .....	248
CharacterSet .....	249
Contents .....	249
See Also .....	249
CloudwatchLogsExportConfiguration .....	250
Contents .....	250
See Also .....	250
ClusterPendingModifiedValues .....	251
Contents .....	251
See Also .....	252
DBCluster .....	254
Contents .....	254
See Also .....	262
DBClusterEndpoint .....	264
Contents .....	264

See Also .....	266
DBClusterMember .....	267
Contents .....	267
See Also .....	267
DBClusterOptionGroupStatus .....	269
Contents .....	269
See Also .....	269
DBClusterParameterGroup .....	270
Contents .....	270
See Also .....	270
DBClusterRole .....	272
Contents .....	272
See Also .....	272
DBClusterSnapshot .....	274
Contents .....	274
See Also .....	278
DBClusterSnapshotAttribute .....	279
Contents .....	279
See Also .....	279
DBClusterSnapshotAttributesResult .....	281
Contents .....	281
See Also .....	281
DBEngineVersion .....	282
Contents .....	282
See Also .....	284
DBInstance .....	285
Contents .....	285
See Also .....	294
DBInstanceStatusInfo .....	295
Contents .....	295
See Also .....	295
DBParameterGroup .....	297
Contents .....	297
See Also .....	297
DBParameterGroupStatus .....	299
Contents .....	299

See Also .....	299
DBSecurityGroupMembership .....	300
Contents .....	300
See Also .....	300
DBSubnetGroup .....	301
Contents .....	301
See Also .....	302
DomainMembership .....	303
Contents .....	303
See Also .....	303
DoubleRange .....	305
Contents .....	305
See Also .....	305
Endpoint .....	306
Contents .....	306
See Also .....	306
EngineDefaults .....	308
Contents .....	308
See Also .....	308
Event .....	310
Contents .....	310
See Also .....	311
EventCategoriesMap .....	312
Contents .....	312
See Also .....	312
EventSubscription .....	313
Contents .....	313
See Also .....	315
Filter .....	316
Contents .....	316
See Also .....	316
GlobalCluster .....	317
Contents .....	317
See Also .....	318
GlobalClusterMember .....	320
Contents .....	320

See Also .....	320
OptionGroupMembership .....	322
Contents .....	322
See Also .....	322
OrderableDBInstanceOption .....	323
Contents .....	323
See Also .....	326
Parameter .....	328
Contents .....	328
See Also .....	329
PendingCloudwatchLogsExports .....	331
Contents .....	331
See Also .....	331
PendingMaintenanceAction .....	332
Contents .....	332
See Also .....	333
PendingModifiedValues .....	334
Contents .....	334
See Also .....	336
Range .....	337
Contents .....	337
See Also .....	337
ResourcePendingMaintenanceActions .....	338
Contents .....	338
See Also .....	338
ServerlessV2ScalingConfiguration .....	339
Contents .....	339
See Also .....	339
ServerlessV2ScalingConfigurationInfo .....	340
Contents .....	340
See Also .....	340
Subnet .....	341
Contents .....	341
See Also .....	341
Tag .....	342
Contents .....	342

---

See Also .....	342
Timezone .....	343
Contents .....	343
See Also .....	343
UpgradeTarget .....	344
Contents .....	344
See Also .....	345
ValidDBInstanceModificationsMessage .....	346
Contents .....	346
See Also .....	346
ValidStorageOptions .....	347
Contents .....	347
See Also .....	347
VpcSecurityGroupMembership .....	349
Contents .....	349
See Also .....	349
<b>Common Parameters .....</b>	<b>350</b>
<b>Common Errors .....</b>	<b>353</b>



# Welcome

Amazon Neptune is a fast, reliable, fully-managed graph database service that makes it easy to build and run applications that work with highly connected datasets. The core of Amazon Neptune is a purpose-built, high-performance graph database engine optimized for storing billions of relationships and querying the graph with milliseconds latency. Amazon Neptune supports popular graph models Property Graph and W3C's RDF, and their respective query languages Apache TinkerPop Gremlin and SPARQL, allowing you to easily build queries that efficiently navigate highly connected datasets. Neptune powers graph use cases such as recommendation engines, fraud detection, knowledge graphs, drug discovery, and network security.

This interface reference for Amazon Neptune contains documentation for a programming or command line interface you can use to manage Amazon Neptune. Note that Amazon Neptune is asynchronous, which means that some interfaces might require techniques such as polling or callback functions to determine when a command has been applied. In this reference, the parameter descriptions indicate whether a command is applied immediately, on the next instance reboot, or during the maintenance window. The reference structure is as follows, and we list following some related topics from the user guide.

This document was last published on July 2, 2024.

# Actions

The following actions are supported:

- [AddRoleToDBCluster](#)
- [AddSourceIdentifierToSubscription](#)
- [AddTagsToResource](#)
- [ApplyPendingMaintenanceAction](#)
- [CopyDBClusterParameterGroup](#)
- [CopyDBClusterSnapshot](#)
- [CopyDBParameterGroup](#)
- [CreateDBCluster](#)
- [CreateDBClusterEndpoint](#)
- [CreateDBClusterParameterGroup](#)
- [CreateDBClusterSnapshot](#)
- [CreateDBInstance](#)
- [CreateDBParameterGroup](#)
- [CreateDBSubnetGroup](#)
- [CreateEventSubscription](#)
- [CreateGlobalCluster](#)
- [DeleteDBCluster](#)
- [DeleteDBClusterEndpoint](#)
- [DeleteDBClusterParameterGroup](#)
- [DeleteDBClusterSnapshot](#)
- [DeleteDBInstance](#)
- [DeleteDBParameterGroup](#)
- [DeleteDBSubnetGroup](#)
- [DeleteEventSubscription](#)
- [DeleteGlobalCluster](#)
- [DescribeDBClusterEndpoints](#)
- [DescribeDBClusterParameterGroups](#)

- [DescribeDBClusterParameters](#)
- [DescribeDBClusters](#)
- [DescribeDBClusterSnapshotAttributes](#)
- [DescribeDBClusterSnapshots](#)
- [DescribeDBEngineVersions](#)
- [DescribeDBInstances](#)
- [DescribeDBParameterGroups](#)
- [DescribeDBParameters](#)
- [DescribeDBSubnetGroups](#)
- [DescribeEngineDefaultClusterParameters](#)
- [DescribeEngineDefaultParameters](#)
- [DescribeEventCategories](#)
- [DescribeEvents](#)
- [DescribeEventSubscriptions](#)
- [DescribeGlobalClusters](#)
- [DescribeOrderableDBInstanceOptions](#)
- [DescribePendingMaintenanceActions](#)
- [DescribeValidDBInstanceModifications](#)
- [FailoverDBCluster](#)
- [FailoverGlobalCluster](#)
- [ListTagsForResource](#)
- [ModifyDBCluster](#)
- [ModifyDBClusterEndpoint](#)
- [ModifyDBClusterParameterGroup](#)
- [ModifyDBClusterSnapshotAttribute](#)
- [ModifyDBInstance](#)
- [ModifyDBParameterGroup](#)
- [ModifyDBSubnetGroup](#)
- [ModifyEventSubscription](#)
- [ModifyGlobalCluster](#)

- [PromoteReadReplicaDBCluster](#)
- [RebootDBInstance](#)
- [RemoveFromGlobalCluster](#)
- [RemoveRoleFromDBCluster](#)
- [RemoveSourceIdentifierFromSubscription](#)
- [RemoveTagsFromResource](#)
- [ResetDBClusterParameterGroup](#)
- [ResetDBParameterGroup](#)
- [RestoreDBClusterFromSnapshot](#)
- [RestoreDBClusterToPointInTime](#)
- [StartDBCluster](#)
- [StopDBCluster](#)

# AddRoleToDBCluster

Associates an Identity and Access Management (IAM) role with an Neptune DB cluster.

## Request Parameters

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

### DBClusterIdentifier

The name of the DB cluster to associate the IAM role with.

Type: String

Required: Yes

### FeatureName

The name of the feature for the Neptune DB cluster that the IAM role is to be associated with. For the list of supported feature names, see [DBEngineVersion](#).

Type: String

Required: No

### RoleArn

The Amazon Resource Name (ARN) of the IAM role to associate with the Neptune DB cluster, for example `arn:aws:iam::123456789012:role/NeptuneAccessRole`.

Type: String

Required: Yes

## Errors

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

### DBClusterNotFoundFault

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

## **DBClusterRoleAlreadyExists**

The specified IAM role Amazon Resource Name (ARN) is already associated with the specified DB cluster.

HTTP Status Code: 400

## **DBClusterRoleQuotaExceeded**

You have exceeded the maximum number of IAM roles that can be associated with the specified DB cluster.

HTTP Status Code: 400

## **InvalidDBClusterStateFault**

The DB cluster is not in a valid state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# AddSourceIdentifierToSubscription

Adds a source identifier to an existing event notification subscription.

## Request Parameters

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

### SourceIdentifier

The identifier of the event source to be added.

Constraints:

- If the source type is a DB instance, then a `DBInstanceIdentifier` must be supplied.
- If the source type is a DB security group, a `DBSecurityGroupName` must be supplied.
- If the source type is a DB parameter group, a `DBParameterGroupName` must be supplied.
- If the source type is a DB snapshot, a `DBSnapshotIdentifier` must be supplied.

Type: String

Required: Yes

### SubscriptionName

The name of the event notification subscription you want to add a source identifier to.

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### EventSubscription

Contains the results of a successful invocation of the [DescribeEventSubscriptions](#) action.

Type: [EventSubscription](#) object

## Errors

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

### SourceNotFound

The source could not be found.

HTTP Status Code: 404

### SubscriptionNotFound

The designated subscription 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# AddTagsToResource

Adds metadata tags to an Amazon Neptune resource. These tags can also be used with cost allocation reporting to track cost associated with Amazon Neptune resources, or used in a Condition statement in an IAM policy for Amazon Neptune.

## Request Parameters

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

### ResourceName

The Amazon Neptune resource that the tags are added to. This value is an Amazon Resource Name (ARN). For information about creating an ARN, see [Constructing an Amazon Resource Name \(ARN\)](#).

Type: String

Required: Yes

### Tags.Tag.N

The tags to be assigned to the Amazon Neptune resource.

Type: Array of [Tag](#) objects

Required: Yes

## Errors

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

### DBClusterNotFoundFault

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### DBInstanceNotFound

*DBInstanceIdentifier* does not refer to an existing DB instance.

HTTP Status Code: 404

## DBSnapshotNotFound

*DBSnapshotIdentifier* does not refer to an existing DB snapshot.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ApplyPendingMaintenanceAction

Applies a pending maintenance action to a resource (for example, to a DB instance).

## Request Parameters

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

### ApplyAction

The pending maintenance action to apply to this resource.

Valid values: `system-update`, `db-upgrade`

Type: String

Required: Yes

### OptInType

A value that specifies the type of opt-in request, or undoes an opt-in request. An opt-in request of type `immediate` can't be undone.

Valid values:

- `immediate` - Apply the maintenance action immediately.
- `next-maintenance` - Apply the maintenance action during the next maintenance window for the resource.
- `undo-opt-in` - Cancel any existing `next-maintenance` opt-in requests.

Type: String

Required: Yes

### ResourceIdentifier

The Amazon Resource Name (ARN) of the resource that the pending maintenance action applies to. For information about creating an ARN, see [Constructing an Amazon Resource Name \(ARN\)](#).

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### ResourcePendingMaintenanceActions

Describes the pending maintenance actions for a resource.

Type: [ResourcePendingMaintenanceActions](#) object

## Errors

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

### ResourceNotFoundFault

The specified resource ID was not 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CopyDBClusterParameterGroup

Copies the specified DB cluster parameter group.

## Request Parameters

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

### SourceDBClusterParameterGroupIdentifier

The identifier or Amazon Resource Name (ARN) for the source DB cluster parameter group. For information about creating an ARN, see [Constructing an Amazon Resource Name \(ARN\)](#).

Constraints:

- Must specify a valid DB cluster parameter group.
- If the source DB cluster parameter group is in the same Amazon Region as the copy, specify a valid DB parameter group identifier, for example `my-db-cluster-param-group`, or a valid ARN.
- If the source DB parameter group is in a different Amazon Region than the copy, specify a valid DB cluster parameter group ARN, for example `arn:aws:rds:us-east-1:123456789012:cluster-pg:custom-cluster-group1`.

Type: String

Required: Yes

### Tags.Tag.N

The tags to be assigned to the copied DB cluster parameter group.

Type: Array of [Tag](#) objects

Required: No

### TargetDBClusterParameterGroupDescription

A description for the copied DB cluster parameter group.

Type: String

Required: Yes

## TargetDBClusterParameterGroupIdentifier

The identifier for the copied DB cluster parameter group.

Constraints:

- Cannot be null, empty, or blank
- Must contain from 1 to 255 letters, numbers, or hyphens
- First character must be a letter
- Cannot end with a hyphen or contain two consecutive hyphens

Example: `my-cluster-param-group1`

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### DBClusterParameterGroup

Contains the details of an Amazon Neptune DB cluster parameter group.

This data type is used as a response element in the [DescribeDBClusterParameterGroups](#) action.

Type: [DBClusterParameterGroup](#) object

## Errors

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

### DBParameterGroupAlreadyExists

A DB parameter group with the same name exists.

HTTP Status Code: 400

### DBParameterGroupNotFound

*DBParameterGroupName* does not refer to an existing DB parameter group.

HTTP Status Code: 404

## **DBParameterGroupQuotaExceeded**

Request would result in user exceeding the allowed number of DB parameter groups.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CopyDBClusterSnapshot

Copies a snapshot of a DB cluster.

To copy a DB cluster snapshot from a shared manual DB cluster snapshot, `SourceDBClusterSnapshotIdentifier` must be the Amazon Resource Name (ARN) of the shared DB cluster snapshot.

## Request Parameters

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

### CopyTags

True to copy all tags from the source DB cluster snapshot to the target DB cluster snapshot, and otherwise false. The default is false.

Type: Boolean

Required: No

### KmsKeyId

The Amazon Amazon KMS key ID for an encrypted DB cluster snapshot. The KMS key ID is the Amazon Resource Name (ARN), KMS key identifier, or the KMS key alias for the KMS encryption key.

If you copy an encrypted DB cluster snapshot from your Amazon account, you can specify a value for `KmsKeyId` to encrypt the copy with a new KMS encryption key. If you don't specify a value for `KmsKeyId`, then the copy of the DB cluster snapshot is encrypted with the same KMS key as the source DB cluster snapshot.

If you copy an encrypted DB cluster snapshot that is shared from another Amazon account, then you must specify a value for `KmsKeyId`.

KMS encryption keys are specific to the Amazon Region that they are created in, and you can't use encryption keys from one Amazon Region in another Amazon Region.

You cannot encrypt an unencrypted DB cluster snapshot when you copy it. If you try to copy an unencrypted DB cluster snapshot and specify a value for the `KmsKeyId` parameter, an error is returned.



Type: String

Required: No

### **PreSignedUrl**

Not currently supported.

Type: String

Required: No

### **SourceDBClusterSnapshotIdentifier**

The identifier of the DB cluster snapshot to copy. This parameter is not case-sensitive.

Constraints:

- Must specify a valid system snapshot in the "available" state.
- Specify a valid DB snapshot identifier.

Example: `my-cluster-snapshot1`

Type: String

Required: Yes

### **Tags.Tag.N**

The tags to assign to the new DB cluster snapshot copy.

Type: Array of [Tag](#) objects

Required: No

### **TargetDBClusterSnapshotIdentifier**

The identifier of the new DB cluster snapshot to create from the source DB cluster snapshot. This parameter is not case-sensitive.

Constraints:

- Must contain from 1 to 63 letters, numbers, or hyphens.
- First character must be a letter.
- Cannot end with a hyphen or contain two consecutive hyphens.

Example: `my-cluster-snapshot2`

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### **DBClusterSnapshot**

Contains the details for an Amazon Neptune DB cluster snapshot

This data type is used as a response element in the [DescribeDBClusterSnapshots](#) action.

Type: [DBClusterSnapshot](#) object

## Errors

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

### **DBClusterSnapshotAlreadyExistsFault**

User already has a DB cluster snapshot with the given identifier.

HTTP Status Code: 400

### **DBClusterSnapshotNotFoundFault**

*DBClusterSnapshotIdentifier* does not refer to an existing DB cluster snapshot.

HTTP Status Code: 404

### **InvalidDBClusterSnapshotStateFault**

The supplied value is not a valid DB cluster snapshot state.

HTTP Status Code: 400

### **InvalidDBClusterStateFault**

The DB cluster is not in a valid state.

HTTP Status Code: 400

### **KMSKeyNotAccessibleFault**

Error accessing KMS key.

HTTP Status Code: 400

### **SnapshotQuotaExceeded**

Request would result in user exceeding the allowed number of DB snapshots.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CopyDBParameterGroup

Copies the specified DB parameter group.

## Request Parameters

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

### SourceDBParameterGroupIdentifier

The identifier or ARN for the source DB parameter group. For information about creating an ARN, see [Constructing an Amazon Resource Name \(ARN\)](#).

Constraints:

- Must specify a valid DB parameter group.
- Must specify a valid DB parameter group identifier, for example `my-db-param-group`, or a valid ARN.

Type: String

Required: Yes

### Tags.Tag.N

The tags to be assigned to the copied DB parameter group.

Type: Array of [Tag](#) objects

Required: No

### TargetDBParameterGroupDescription

A description for the copied DB parameter group.

Type: String

Required: Yes

### TargetDBParameterGroupIdentifier

The identifier for the copied DB parameter group.

Constraints:

- Cannot be null, empty, or blank.
- Must contain from 1 to 255 letters, numbers, or hyphens.
- First character must be a letter.
- Cannot end with a hyphen or contain two consecutive hyphens.

Example: `my-db-parameter-group`

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### **DBParameterGroup**

Contains the details of an Amazon Neptune DB parameter group.

This data type is used as a response element in the [DescribeDBParameterGroups](#) action.

Type: [DBParameterGroup](#) object

## Errors

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

### **DBParameterGroupAlreadyExists**

A DB parameter group with the same name exists.

HTTP Status Code: 400

### **DBParameterGroupNotFound**

*DBParameterGroupName* does not refer to an existing DB parameter group.

HTTP Status Code: 404

### **DBParameterGroupQuotaExceeded**

Request would result in user exceeding the allowed number of DB parameter groups.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateDBCluster

Creates a new Amazon Neptune DB cluster.

You can use the `ReplicationSourceIdentifier` parameter to create the DB cluster as a Read Replica of another DB cluster or Amazon Neptune DB instance.

Note that when you create a new cluster using `CreateDBCluster` directly, deletion protection is disabled by default (when you create a new production cluster in the console, deletion protection is enabled by default). You can only delete a DB cluster if its `DeletionProtection` field is set to `false`.

## Request Parameters

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

### AvailabilityZones.AvailabilityZone.N

A list of EC2 Availability Zones that instances in the DB cluster can be created in.

Type: Array of strings

Required: No

### BackupRetentionPeriod

The number of days for which automated backups are retained. You must specify a minimum value of 1.

Default: 1

Constraints:

- Must be a value from 1 to 35

Type: Integer

Required: No

### CharacterSetName

*(Not supported by Neptune)*

Type: String

Required: No

### **CopyTagsToSnapshot**

*If set to true, tags are copied to any snapshot of the DB cluster that is created.*

Type: Boolean

Required: No

### **DatabaseName**

The name for your database of up to 64 alpha-numeric characters. If you do not provide a name, Amazon Neptune will not create a database in the DB cluster you are creating.

Type: String

Required: No

### **DBClusterIdentifier**

The DB cluster identifier. This parameter is stored as a lowercase string.

Constraints:

- Must contain from 1 to 63 letters, numbers, or hyphens.
- First character must be a letter.
- Cannot end with a hyphen or contain two consecutive hyphens.

Example: `my-cluster1`

Type: String

Required: Yes

### **DBClusterParameterGroupName**

The name of the DB cluster parameter group to associate with this DB cluster. If this argument is omitted, the default is used.

Constraints:

- If supplied, must match the name of an existing `DBClusterParameterGroup`.

Type: String

Required: No



## **DBSubnetGroupName**

A DB subnet group to associate with this DB cluster.

Constraints: Must match the name of an existing DBSubnetGroup. Must not be default.

Example: mySubnetgroup

Type: String

Required: No

## **DeletionProtection**

A value that indicates whether the DB cluster has deletion protection enabled. The database can't be deleted when deletion protection is enabled. By default, deletion protection is enabled.

Type: Boolean

Required: No

## **EnableCloudwatchLogsExports.member.N**

A list of the log types that this DB cluster should export to CloudWatch Logs. Valid log types are: `audit` (to publish audit logs) and `slowquery` (to publish slow-query logs). See [Publishing Neptune logs to Amazon CloudWatch logs](#).

Type: Array of strings

Required: No

## **EnableIAMDatabaseAuthentication**

If set to `true`, enables Amazon Identity and Access Management (IAM) authentication for the entire DB cluster (this cannot be set at an instance level).

Default: `false`.

Type: Boolean

Required: No

## **Engine**

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

Valid Values: neptune

Type: String

Required: Yes

### **EngineVersion**

The version number of the database engine to use for the new DB cluster.

Example: 1.2.1.0

Type: String

Required: No

### **GlobalClusterIdentifier**

The ID of the Neptune global database to which this new DB cluster should be added.

Type: String

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

Pattern: [A-Za-z][0-9A-Za-z-:.\_]\*

Required: No

### **KmsKeyId**

The Amazon KMS key identifier for an encrypted DB cluster.

The KMS key identifier is the Amazon Resource Name (ARN) for the KMS encryption key. If you are creating a DB cluster with the same Amazon account that owns the KMS encryption key used to encrypt the new DB cluster, then you can use the KMS key alias instead of the ARN for the KMS encryption key.

If an encryption key is not specified in `KmsKeyId`:

- If `ReplicationSourceIdentifier` identifies an encrypted source, then Amazon Neptune will use the encryption key used to encrypt the source. Otherwise, Amazon Neptune will use your default encryption key.
- If the `StorageEncrypted` parameter is true and `ReplicationSourceIdentifier` is not specified, then Amazon Neptune will use your default encryption key.

Amazon KMS creates the default encryption key for your Amazon account. Your Amazon account has a different default encryption key for each Amazon Region.

If you create a Read Replica of an encrypted DB cluster in another Amazon Region, you must set `KmsKeyId` to a KMS key ID that is valid in the destination Amazon Region. This key is used to encrypt the Read Replica in that Amazon Region.

Type: String

Required: No

### **MasterUsername**

Not supported by Neptune.

Type: String

Required: No

### **MasterUserPassword**

Not supported by Neptune.

Type: String

Required: No

### **OptionGroupName**

*(Not supported by Neptune)*

Type: String

Required: No

### **Port**

The port number on which the instances in the DB cluster accept connections.

Default: 8182

Type: Integer

Required: No

## PreferredBackupWindow

The daily time range during which automated backups are created if automated backups are enabled using the BackupRetentionPeriod parameter.

The default is a 30-minute window selected at random from an 8-hour block of time for each Amazon Region. To see the time blocks available, see [Neptune Maintenance Window](#) in the *Amazon Neptune User Guide*.

Constraints:

- Must be in the format hh24:mi-hh24:mi.
- Must be in Universal Coordinated Time (UTC).
- Must not conflict with the preferred maintenance window.
- Must be at least 30 minutes.

Type: String

Required: No

## PreferredMaintenanceWindow

The weekly time range during which system maintenance can occur, in Universal Coordinated Time (UTC).

Format: ddd:hh24:mi-ddd:hh24:mi

The default is a 30-minute window selected at random from an 8-hour block of time for each Amazon Region, occurring on a random day of the week. To see the time blocks available, see [Neptune Maintenance Window](#) in the *Amazon Neptune User Guide*.

Valid Days: Mon, Tue, Wed, Thu, Fri, Sat, Sun.

Constraints: Minimum 30-minute window.

Type: String

Required: No

## PreSignedUrl

This parameter is not currently supported.

Type: String

Required: No

### ReplicationSourceIdentifier

The Amazon Resource Name (ARN) of the source DB instance or DB cluster if this DB cluster is created as a Read Replica.

Type: String

Required: No

### ServerlessV2ScalingConfiguration

Contains the scaling configuration of a Neptune Serverless DB cluster.

For more information, see [Using Amazon Neptune Serverless](#) in the *Amazon Neptune User Guide*.

Type: [ServerlessV2ScalingConfiguration](#) object

Required: No

### StorageEncrypted

Specifies whether the DB cluster is encrypted.

Type: Boolean

Required: No

### StorageType

The storage type for the new DB cluster.

Valid Values:

- **standard** – ( *the default* ) Configures cost-effective database storage for applications with moderate to small I/O usage. When set to standard, the storage type is not returned in the response.
- **iopt1** – Enables [I/O-Optimized storage](#) that's designed to meet the needs of I/O-intensive graph workloads that require predictable pricing with low I/O latency and consistent I/O throughput.

Neptune I/O-Optimized storage is only available starting with engine release 1.3.0.0.

Type: String

Required: No

### **Tags.Tag.N**

The tags to assign to the new DB cluster.

Type: Array of [Tag](#) objects

Required: No

### **VpcSecurityGroupIds.VpcSecurityGroupId.N**

A list of EC2 VPC security groups to associate with this DB cluster.

Type: Array of strings

Required: No

## **Response Elements**

The following element is returned by the service.

### **DBCluster**

Contains the details of an Amazon Neptune DB cluster.

This data type is used as a response element in the [DescribeDBClusters](#).

Type: [DBCluster](#) object

## **Errors**

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

### **DBClusterAlreadyExistsFault**

User already has a DB cluster with the given identifier.

HTTP Status Code: 400

### **DBClusterNotFoundFault**

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### **DBClusterParameterGroupNotFound**

*DBClusterParameterGroupName* does not refer to an existing DB Cluster parameter group.

HTTP Status Code: 404

### **DBClusterQuotaExceededFault**

User attempted to create a new DB cluster and the user has already reached the maximum allowed DB cluster quota.

HTTP Status Code: 403

### **DBInstanceNotFound**

*DBInstanceIdentifier* does not refer to an existing DB instance.

HTTP Status Code: 404

### **DBSubnetGroupDoesNotCoverEnoughAZs**

Subnets in the DB subnet group should cover at least two Availability Zones unless there is only one Availability Zone.

HTTP Status Code: 400

### **DBSubnetGroupNotFoundFault**

*DBSubnetGroupName* does not refer to an existing DB subnet group.

HTTP Status Code: 404

### **GlobalClusterNotFoundFault**

The `GlobalClusterIdentifier` doesn't refer to an existing global database cluster.

HTTP Status Code: 404

### **InsufficientStorageClusterCapacity**

There is insufficient storage available for the current action. You may be able to resolve this error by updating your subnet group to use different Availability Zones that have more storage available.

HTTP Status Code: 400

**InvalidDBClusterStateFault**

The DB cluster is not in a valid state.

HTTP Status Code: 400

**InvalidDBInstanceState**

The specified DB instance is not in the *available* state.

HTTP Status Code: 400

**InvalidDBSubnetGroupStateFault**

The DB subnet group cannot be deleted because it is in use.

HTTP Status Code: 400

**InvalidGlobalClusterStateFault**

The global cluster is in an invalid state and can't perform the requested operation.

HTTP Status Code: 400

**InvalidSubnet**

The requested subnet is invalid, or multiple subnets were requested that are not all in a common VPC.

HTTP Status Code: 400

**InvalidVPCNetworkStateFault**

DB subnet group does not cover all Availability Zones after it is created because users' change.

HTTP Status Code: 400

**KMSKeyNotAccessibleFault**

Error accessing KMS key.

HTTP Status Code: 400

**StorageQuotaExceeded**

Request would result in user exceeding the allowed amount of storage available across all DB instances.



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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateDBClusterEndpoint

Creates a new custom endpoint and associates it with an Amazon Neptune DB cluster.

## Request Parameters

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

### DBClusterEndpointIdentifier

The identifier to use for the new endpoint. This parameter is stored as a lowercase string.

Type: String

Required: Yes

### DBClusterIdentifier

The DB cluster identifier of the DB cluster associated with the endpoint. This parameter is stored as a lowercase string.

Type: String

Required: Yes

### EndpointType

The type of the endpoint. One of: `READER`, `WRITER`, `ANY`.

Type: String

Required: Yes

### ExcludedMembers.member.N

List of DB instance identifiers that aren't part of the custom endpoint group. All other eligible instances are reachable through the custom endpoint. Only relevant if the list of static members is empty.

Type: Array of strings

Required: No

### StaticMembers.member.N

List of DB instance identifiers that are part of the custom endpoint group.

Type: Array of strings

Required: No

### **Tags.Tag.N**

The tags to be assigned to the Amazon Neptune resource.

Type: Array of [Tag](#) objects

Required: No

## **Response Elements**

The following elements are returned by the service.

### **CustomEndpointType**

The type associated with a custom endpoint. One of: `READER`, `WRITER`, `ANY`.

Type: String

### **DBClusterEndpointArn**

The Amazon Resource Name (ARN) for the endpoint.

Type: String

### **DBClusterEndpointIdentifier**

The identifier associated with the endpoint. This parameter is stored as a lowercase string.

Type: String

### **DBClusterEndpointResourceIdentifier**

A unique system-generated identifier for an endpoint. It remains the same for the whole life of the endpoint.

Type: String

### **DBClusterIdentifier**

The DB cluster identifier of the DB cluster associated with the endpoint. This parameter is stored as a lowercase string.

Type: String

### **Endpoint**

The DNS address of the endpoint.

Type: String

### **EndpointType**

The type of the endpoint. One of: `READER`, `WRITER`, `CUSTOM`.

Type: String

### **ExcludedMembers.member.N**

List of DB instance identifiers that aren't part of the custom endpoint group. All other eligible instances are reachable through the custom endpoint. Only relevant if the list of static members is empty.

Type: Array of strings

### **StaticMembers.member.N**

List of DB instance identifiers that are part of the custom endpoint group.

Type: Array of strings

### **Status**

The current status of the endpoint. One of: `creating`, `available`, `deleting`, `inactive`, `modifying`. The `inactive` state applies to an endpoint that cannot be used for a certain kind of cluster, such as a `writer` endpoint for a read-only secondary cluster in a global database.

Type: String

## **Errors**

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

### **DBClusterEndpointAlreadyExistsFault**

The specified custom endpoint cannot be created because it already exists.

HTTP Status Code: 400

## DBClusterEndpointQuotaExceededFault

The cluster already has the maximum number of custom endpoints.

HTTP Status Code: 403

## DBClusterNotFoundFault

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

## DBInstanceNotFound

*DBInstanceIdentifier* does not refer to an existing DB instance.

HTTP Status Code: 404

## InvalidDBClusterStateFault

The DB cluster is not in a valid state.

HTTP Status Code: 400

## InvalidDBInstanceState

The specified DB instance is not in the *available* state.

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 PHP V3](#)

- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateDBClusterParameterGroup

Creates a new DB cluster parameter group.

Parameters in a DB cluster parameter group apply to all of the instances in a DB cluster.

A DB cluster parameter group is initially created with the default parameters for the database engine used by instances in the DB cluster. To provide custom values for any of the parameters, you must modify the group after creating it using [ModifyDBClusterParameterGroup](#). Once you've created a DB cluster parameter group, you need to associate it with your DB cluster using [ModifyDBCluster](#). When you associate a new DB cluster parameter group with a running DB cluster, you need to reboot the DB instances in the DB cluster without failover for the new DB cluster parameter group and associated settings to take effect.

## Important

After you create a DB cluster parameter group, you should wait at least 5 minutes before creating your first DB cluster that uses that DB cluster parameter group as the default parameter group. This allows Amazon Neptune to fully complete the create action before the DB cluster parameter group is used as the default for a new DB cluster. This is especially important for parameters that are critical when creating the default database for a DB cluster, such as the character set for the default database defined by the `character_set_database` parameter. You can use the *Parameter Groups* option of the [Amazon Neptune console](#) or the [DescribeDBClusterParameters](#) command to verify that your DB cluster parameter group has been created or modified.

## Request Parameters

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

### DBClusterParameterGroupName

The name of the DB cluster parameter group.

Constraints:

- Must match the name of an existing `DBClusterParameterGroup`.

**Note**

This value is stored as a lowercase string.

Type: String

Required: Yes

**DBParameterGroupFamily**

The DB cluster parameter group family name. A DB cluster parameter group can be associated with one and only one DB cluster parameter group family, and can be applied only to a DB cluster running a database engine and engine version compatible with that DB cluster parameter group family.

Type: String

Required: Yes

**Description**

The description for the DB cluster parameter group.

Type: String

Required: Yes

**Tags.Tag.N**

The tags to be assigned to the new DB cluster parameter group.

Type: Array of [Tag](#) objects

Required: No

**Response Elements**

The following element is returned by the service.

**DBClusterParameterGroup**

Contains the details of an Amazon Neptune DB cluster parameter group.



This data type is used as a response element in the [DescribeDBClusterParameterGroups](#) action.

Type: [DBClusterParameterGroup](#) object

## Errors

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

### **DBParameterGroupAlreadyExists**

A DB parameter group with the same name exists.

HTTP Status Code: 400

### **DBParameterGroupQuotaExceeded**

Request would result in user exceeding the allowed number of DB parameter groups.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateDBClusterSnapshot

Creates a snapshot of a DB cluster.

## Request Parameters

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

### DBClusterIdentifier

The identifier of the DB cluster to create a snapshot for. This parameter is not case-sensitive.

Constraints:

- Must match the identifier of an existing DBCluster.

Example: `my-cluster1`

Type: String

Required: Yes

### DBClusterSnapshotIdentifier

The identifier of the DB cluster snapshot. This parameter is stored as a lowercase string.

Constraints:

- Must contain from 1 to 63 letters, numbers, or hyphens.
- First character must be a letter.
- Cannot end with a hyphen or contain two consecutive hyphens.

Example: `my-cluster1-snapshot1`

Type: String

Required: Yes

### Tags.Tag.N

The tags to be assigned to the DB cluster snapshot.

Type: Array of [Tag](#) objects

Required: No

## Response Elements

The following element is returned by the service.

### **DBClusterSnapshot**

Contains the details for an Amazon Neptune DB cluster snapshot

This data type is used as a response element in the [DescribeDBClusterSnapshots](#) action.

Type: [DBClusterSnapshot](#) object

## Errors

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

### **DBClusterNotFoundFault**

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### **DBClusterSnapshotAlreadyExistsFault**

User already has a DB cluster snapshot with the given identifier.

HTTP Status Code: 400

### **InvalidDBClusterSnapshotStateFault**

The supplied value is not a valid DB cluster snapshot state.

HTTP Status Code: 400

### **InvalidDBClusterStateFault**

The DB cluster is not in a valid state.

HTTP Status Code: 400

### **SnapshotQuotaExceeded**

Request would result in user exceeding the allowed number of DB snapshots.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateDBInstance

Creates a new DB instance.

## Request Parameters

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

### AllocatedStorage

Not supported by Neptune.

Type: Integer

Required: No

### AutoMinorVersionUpgrade

Indicates that minor engine upgrades are applied automatically to the DB instance during the maintenance window.

Default: `true`

Type: Boolean

Required: No

### AvailabilityZone

The EC2 Availability Zone that the DB instance is created in

Default: A random, system-chosen Availability Zone in the endpoint's Amazon Region.

Example: `us-east-1d`

Constraint: The `AvailabilityZone` parameter can't be specified if the `MultiAZ` parameter is set to `true`. The specified Availability Zone must be in the same Amazon Region as the current endpoint.

Type: String

Required: No

## BackupRetentionPeriod

The number of days for which automated backups are retained.

Not applicable. The retention period for automated backups is managed by the DB cluster. For more information, see [CreateDBCluster](#).

Default: 1

Constraints:

- Must be a value from 0 to 35
- Cannot be set to 0 if the DB instance is a source to Read Replicas

Type: Integer

Required: No

## CharacterSetName

*(Not supported by Neptune)*

Type: String

Required: No

## CopyTagsToSnapshot

True to copy all tags from the DB instance to snapshots of the DB instance, and otherwise false. The default is false.

Type: Boolean

Required: No

## DBClusterIdentifier

The identifier of the DB cluster that the instance will belong to.

For information on creating a DB cluster, see [CreateDBCluster](#).

Type: String

Type: String

Required: Yes

### **DBInstanceClass**

The compute and memory capacity of the DB instance, for example, `db.m4.large`. Not all DB instance classes are available in all Amazon Regions.

Type: String

Required: Yes

### **DBInstanceIdentifier**

The DB instance identifier. This parameter is stored as a lowercase string.

Constraints:

- Must contain from 1 to 63 letters, numbers, or hyphens.
- First character must be a letter.
- Cannot end with a hyphen or contain two consecutive hyphens.

Example: `mydbinstance`

Type: String

Required: Yes

### **DBName**

Not supported.

Type: String

Required: No

### **DBParameterGroupName**

The name of the DB parameter group to associate with this DB instance. If this argument is omitted, the default `DBParameterGroup` for the specified engine is used.

Constraints:

- Must be 1 to 255 letters, numbers, or hyphens.
- First character must be a letter

- Cannot end with a hyphen or contain two consecutive hyphens

Type: String

Required: No

### **DBSecurityGroups.DBSecurityGroupName.N**

A list of DB security groups to associate with this DB instance.

Default: The default DB security group for the database engine.

Type: Array of strings

Required: No

### **DBSubnetGroupName**

A DB subnet group to associate with this DB instance.

If there is no DB subnet group, then it is a non-VPC DB instance.

Type: String

Required: No

### **DeletionProtection**

A value that indicates whether the DB instance has deletion protection enabled. The database can't be deleted when deletion protection is enabled. By default, deletion protection is disabled. See [Deleting a DB Instance](#).

DB instances in a DB cluster can be deleted even when deletion protection is enabled in their parent DB cluster.

Type: Boolean

Required: No

### **Domain**

Specify the Active Directory Domain to create the instance in.

Type: String

Required: No



**DomainIAMRoleName**

Specify the name of the IAM role to be used when making API calls to the Directory Service.

Type: String

Required: No

**EnableCloudwatchLogsExports.member.N**

The list of log types that need to be enabled for exporting to CloudWatch Logs.

Type: Array of strings

Required: No

**EnableIAMDatabaseAuthentication**

Not supported by Neptune (ignored).

Type: Boolean

Required: No

**EnablePerformanceInsights**

*(Not supported by Neptune)*

Type: Boolean

Required: No

**Engine**

The name of the database engine to be used for this instance.

Valid Values: neptune

Type: String

Required: Yes

**EngineVersion**

The version number of the database engine to use. Currently, setting this parameter has no effect.

Type: String

Required: No

### **Iops**

The amount of Provisioned IOPS (input/output operations per second) to be initially allocated for the DB instance.

Type: Integer

Required: No

### **KmsKeyId**

The Amazon KMS key identifier for an encrypted DB instance.

The KMS key identifier is the Amazon Resource Name (ARN) for the KMS encryption key. If you are creating a DB instance with the same Amazon account that owns the KMS encryption key used to encrypt the new DB instance, then you can use the KMS key alias instead of the ARN for the KM encryption key.

Not applicable. The KMS key identifier is managed by the DB cluster. For more information, see [CreateDBCluster](#).

If the `StorageEncrypted` parameter is true, and you do not specify a value for the `KmsKeyId` parameter, then Amazon Neptune will use your default encryption key. Amazon KMS creates the default encryption key for your Amazon account. Your Amazon account has a different default encryption key for each Amazon Region.

Type: String

Required: No

### **LicenseModel**

License model information for this DB instance.

Valid values: `license-included` | `bring-your-own-license` | `general-public-license`

Type: String

Required: No

## MasterUsername

Not supported by Neptune.

Type: String

Required: No

## MasterUserPassword

Not supported by Neptune.

Type: String

Required: No

## MonitoringInterval

The interval, in seconds, between points when Enhanced Monitoring metrics are collected for the DB instance. To disable collecting Enhanced Monitoring metrics, specify 0. The default is 0.

If `MonitoringRoleArn` is specified, then you must also set `MonitoringInterval` to a value other than 0.

Valid Values: 0, 1, 5, 10, 15, 30, 60

Type: Integer

Required: No

## MonitoringRoleArn

The ARN for the IAM role that permits Neptune to send enhanced monitoring metrics to Amazon CloudWatch Logs. For example, `arn:aws:iam:123456789012:role/emaccess`.

If `MonitoringInterval` is set to a value other than 0, then you must supply a `MonitoringRoleArn` value.

Type: String

Required: No

## MultiAZ

Specifies if the DB instance is a Multi-AZ deployment. You can't set the `AvailabilityZone` parameter if the `MultiAZ` parameter is set to true.

Type: Boolean

Required: No

### **OptionGroupName**

*(Not supported by Neptune)*

Type: String

Required: No

### **PerformanceInsightsKMSKeyId**

*(Not supported by Neptune)*

Type: String

Required: No

### **Port**

The port number on which the database accepts connections.

Not applicable. The port is managed by the DB cluster. For more information, see [CreateDBCluster](#).

Default: 8182

Type: Integer

Type: Integer

Required: No

### **PreferredBackupWindow**

The daily time range during which automated backups are created.

Not applicable. The daily time range for creating automated backups is managed by the DB cluster. For more information, see [CreateDBCluster](#).

Type: String

Required: No

## PreferredMaintenanceWindow

The time range each week during which system maintenance can occur, in Universal Coordinated Time (UTC).

Format: ddd:hh24:mi-dd:hh24:mi

The default is a 30-minute window selected at random from an 8-hour block of time for each Amazon Region, occurring on a random day of the week.

Valid Days: Mon, Tue, Wed, Thu, Fri, Sat, Sun.

Constraints: Minimum 30-minute window.

Type: String

Required: No

## PromotionTier

A value that specifies the order in which an Read Replica is promoted to the primary instance after a failure of the existing primary instance.

Default: 1

Valid Values: 0 - 15

Type: Integer

Required: No

## PubliclyAccessible

*This parameter has been deprecated.*

This flag should no longer be used.

Type: Boolean

Required: No

## StorageEncrypted

Specifies whether the DB instance is encrypted.

Not applicable. The encryption for DB instances is managed by the DB cluster. For more information, see [CreateDBCluster](#).

Default: false

Type: Boolean

Required: No

### **StorageType**

Not applicable. In Neptune the storage type is managed at the DB Cluster level.

Type: String

Required: No

### **Tags.Tag.N**

The tags to assign to the new instance.

Type: Array of [Tag](#) objects

Required: No

### **TdeCredentialArn**

The ARN from the key store with which to associate the instance for TDE encryption.

Type: String

Required: No

### **TdeCredentialPassword**

The password for the given ARN from the key store in order to access the device.

Type: String

Required: No

### **Timezone**

The time zone of the DB instance.

Type: String

Required: No

### **VpcSecurityGroupIds.VpcSecurityGroupId.N**

A list of EC2 VPC security groups to associate with this DB instance.

Not applicable. The associated list of EC2 VPC security groups is managed by the DB cluster. For more information, see [CreateDBCluster](#).

Default: The default EC2 VPC security group for the DB subnet group's VPC.

Type: Array of strings

Required: No

## **Response Elements**

The following element is returned by the service.

### **DBInstance**

Contains the details of an Amazon Neptune DB instance.

This data type is used as a response element in the [DescribeDBInstances](#) action.

Type: [DBInstance](#) object

## **Errors**

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

### **AuthorizationNotFound**

Specified CIDRIP or EC2 security group is not authorized for the specified DB security group.

Neptune may not also be authorized via IAM to perform necessary actions on your behalf.

HTTP Status Code: 404

### **DBClusterNotFoundFault**

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### **DBInstanceAlreadyExists**

User already has a DB instance with the given identifier.

HTTP Status Code: 400

### **DBParameterGroupNotFound**

*DBParameterGroupName* does not refer to an existing DB parameter group.

HTTP Status Code: 404

### **DBSecurityGroupNotFound**

*DBSecurityGroupName* does not refer to an existing DB security group.

HTTP Status Code: 404

### **DBSubnetGroupDoesNotCoverEnoughAZs**

Subnets in the DB subnet group should cover at least two Availability Zones unless there is only one Availability Zone.

HTTP Status Code: 400

### **DBSubnetGroupNotFoundFault**

*DBSubnetGroupName* does not refer to an existing DB subnet group.

HTTP Status Code: 404

### **DomainNotFoundFault**

*Domain* does not refer to an existing Active Directory Domain.

HTTP Status Code: 404

### **InstanceQuotaExceeded**

Request would result in user exceeding the allowed number of DB instances.

HTTP Status Code: 400

### **InsufficientDBInstanceCapacity**

Specified DB instance class is not available in the specified Availability Zone.



HTTP Status Code: 400

### **InvalidDBClusterStateFault**

The DB cluster is not in a valid state.

HTTP Status Code: 400

### **InvalidSubnet**

The requested subnet is invalid, or multiple subnets were requested that are not all in a common VPC.

HTTP Status Code: 400

### **InvalidVPCNetworkStateFault**

DB subnet group does not cover all Availability Zones after it is created because users' change.

HTTP Status Code: 400

### **KMSKeyNotAccessibleFault**

Error accessing KMS key.

HTTP Status Code: 400

### **OptionGroupNotFoundFault**

The designated option group could not be found.

HTTP Status Code: 404

### **ProvisionedIopsNotAvailableInAZFault**

Provisioned IOPS not available in the specified Availability Zone.

HTTP Status Code: 400

### **StorageQuotaExceeded**

Request would result in user exceeding the allowed amount of storage available across all DB instances.

HTTP Status Code: 400

### **StorageTypeNotSupported**

*StorageType* specified cannot be associated with the DB Instance.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateDBParameterGroup

Creates a new DB parameter group.

A DB parameter group is initially created with the default parameters for the database engine used by the DB instance. To provide custom values for any of the parameters, you must modify the group after creating it using *ModifyDBParameterGroup*. Once you've created a DB parameter group, you need to associate it with your DB instance using *ModifyDBInstance*. When you associate a new DB parameter group with a running DB instance, you need to reboot the DB instance without failover for the new DB parameter group and associated settings to take effect.

## Important

After you create a DB parameter group, you should wait at least 5 minutes before creating your first DB instance that uses that DB parameter group as the default parameter group. This allows Amazon Neptune to fully complete the create action before the parameter group is used as the default for a new DB instance. This is especially important for parameters that are critical when creating the default database for a DB instance, such as the character set for the default database defined by the `character_set_database` parameter. You can use the *Parameter Groups* option of the Amazon Neptune console or the *DescribeDBParameters* command to verify that your DB parameter group has been created or modified.

## Request Parameters

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

### DBParameterGroupFamily

The DB parameter group family name. A DB parameter group can be associated with one and only one DB parameter group family, and can be applied only to a DB instance running a database engine and engine version compatible with that DB parameter group family.

Type: String

Required: Yes

### DBParameterGroupName

The name of the DB parameter group.

**Constraints:**

- Must be 1 to 255 letters, numbers, or hyphens.
- First character must be a letter
- Cannot end with a hyphen or contain two consecutive hyphens

**Note**

This value is stored as a lowercase string.

Type: String

Required: Yes

**Description**

The description for the DB parameter group.

Type: String

Required: Yes

**Tags.Tag.N**

The tags to be assigned to the new DB parameter group.

Type: Array of [Tag](#) objects

Required: No

**Response Elements**

The following element is returned by the service.

**DBParameterGroup**

Contains the details of an Amazon Neptune DB parameter group.

This data type is used as a response element in the [DescribeDBParameterGroups](#) action.

Type: [DBParameterGroup](#) object

## Errors

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

### **DBParameterGroupAlreadyExists**

A DB parameter group with the same name exists.

HTTP Status Code: 400

### **DBParameterGroupQuotaExceeded**

Request would result in user exceeding the allowed number of DB parameter groups.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateDBSubnetGroup

Creates a new DB subnet group. DB subnet groups must contain at least one subnet in at least two AZs in the Amazon Region.

## Request Parameters

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

### DBSubnetGroupDescription

The description for the DB subnet group.

Type: String

Required: Yes

### DBSubnetGroupName

The name for the DB subnet group. This value is stored as a lowercase string.

Constraints: Must contain no more than 255 letters, numbers, periods, underscores, spaces, or hyphens. Must not be default.

Example: mySubnetgroup

Type: String

Required: Yes

### SubnetIds.SubnetIdentifier.N

The EC2 Subnet IDs for the DB subnet group.

Type: Array of strings

Required: Yes

### Tags.Tag.N

The tags to be assigned to the new DB subnet group.

Type: Array of [Tag](#) objects

Required: No

## Response Elements

The following element is returned by the service.

### **DBSubnetGroup**

Contains the details of an Amazon Neptune DB subnet group.

This data type is used as a response element in the [DescribeDBSubnetGroups](#) action.

Type: [DBSubnetGroup](#) object

## Errors

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

### **DBSubnetGroupAlreadyExists**

*DBSubnetGroupName* is already used by an existing DB subnet group.

HTTP Status Code: 400

### **DBSubnetGroupDoesNotCoverEnoughAZs**

Subnets in the DB subnet group should cover at least two Availability Zones unless there is only one Availability Zone.

HTTP Status Code: 400

### **DBSubnetGroupQuotaExceeded**

Request would result in user exceeding the allowed number of DB subnet groups.

HTTP Status Code: 400

### **DBSubnetQuotaExceededFault**

Request would result in user exceeding the allowed number of subnets in a DB subnet groups.

HTTP Status Code: 400

### **InvalidSubnet**

The requested subnet is invalid, or multiple subnets were requested that are not all in a common VPC.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# CreateEventSubscription

Creates an event notification subscription. This action requires a topic ARN (Amazon Resource Name) created by either the Neptune console, the SNS console, or the SNS API. To obtain an ARN with SNS, you must create a topic in Amazon SNS and subscribe to the topic. The ARN is displayed in the SNS console.

You can specify the type of source (`SourceType`) you want to be notified of, provide a list of Neptune sources (`SourceIds`) that triggers the events, and provide a list of event categories (`EventCategories`) for events you want to be notified of. For example, you can specify `SourceType = db-instance`, `SourceIds = mydbinstance1, mydbinstance2` and `EventCategories = Availability, Backup`.

If you specify both the `SourceType` and `SourceIds`, such as `SourceType = db-instance` and `SourceIdentifier = myDBInstance1`, you are notified of all the db-instance events for the specified source. If you specify a `SourceType` but do not specify a `SourceIdentifier`, you receive notice of the events for that source type for all your Neptune sources. If you do not specify either the `SourceType` nor the `SourceIdentifier`, you are notified of events generated from all Neptune sources belonging to your customer account.

## Request Parameters

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

### Enabled

A Boolean value; set to **true** to activate the subscription, set to **false** to create the subscription but not active it.

Type: Boolean

Required: No

### EventCategories.EventCategory.N

A list of event categories for a `SourceType` that you want to subscribe to. You can see a list of the categories for a given `SourceType` by using the **DescribeEventCategories** action.

Type: Array of strings

Required: No

## SnsTopicArn

The Amazon Resource Name (ARN) of the SNS topic created for event notification. The ARN is created by Amazon SNS when you create a topic and subscribe to it.

Type: String

Required: Yes

## SourceIds.SourceId.N

The list of identifiers of the event sources for which events are returned. If not specified, then all sources are included in the response. An identifier must begin with a letter and must contain only ASCII letters, digits, and hyphens; it can't end with a hyphen or contain two consecutive hyphens.

Constraints:

- If `SourceIds` are supplied, `SourceType` must also be provided.
- If the source type is a DB instance, then a `DBInstanceIdentifier` must be supplied.
- If the source type is a DB security group, a `DBSecurityGroupName` must be supplied.
- If the source type is a DB parameter group, a `DBParameterGroupName` must be supplied.
- If the source type is a DB snapshot, a `DBSnapshotIdentifier` must be supplied.

Type: Array of strings

Required: No

## SourceType

The type of source that is generating the events. For example, if you want to be notified of events generated by a DB instance, you would set this parameter to `db-instance`. If this value is not specified, all events are returned.

Valid values: `db-instance` | `db-cluster` | `db-parameter-group` | `db-security-group` | `db-snapshot` | `db-cluster-snapshot`

Type: String

Required: No

## SubscriptionName

The name of the subscription.

Constraints: The name must be less than 255 characters.

Type: String

Required: Yes

### **Tags.Tag.N**

The tags to be applied to the new event subscription.

Type: Array of [Tag](#) objects

Required: No

## **Response Elements**

The following element is returned by the service.

### **EventSubscription**

Contains the results of a successful invocation of the [DescribeEventSubscriptions](#) action.

Type: [EventSubscription](#) object

## **Errors**

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

### **EventSubscriptionQuotaExceeded**

You have exceeded the number of events you can subscribe to.

HTTP Status Code: 400

### **SNSInvalidTopic**

The SNS topic is invalid.

HTTP Status Code: 400

### **SNSNoAuthorization**

There is no SNS authorization.

HTTP Status Code: 400

### **SNSTopicArnNotFound**

The ARN of the SNS topic could not be found.

HTTP Status Code: 404

### **SourceNotFound**

The source could not be found.

HTTP Status Code: 404

### **SubscriptionAlreadyExist**

This subscription already exists.

HTTP Status Code: 400

### **SubscriptionCategoryNotFound**

The designated subscription category 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateGlobalCluster

Creates a Neptune global database spread across multiple Amazon Regions. The global database contains a single primary cluster with read-write capability, and read-only secondary clusters that receive data from the primary cluster through high-speed replication performed by the Neptune storage subsystem.

You can create a global database that is initially empty, and then add a primary cluster and secondary clusters to it, or you can specify an existing Neptune cluster during the create operation to become the primary cluster of the global database.

## Request Parameters

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

### DeletionProtection

The deletion protection setting for the new global database. The global database can't be deleted when deletion protection is enabled.

Type: Boolean

Required: No

### Engine

The name of the database engine to be used in the global database.

Valid values: neptune

Type: String

Required: No

### EngineVersion

The Neptune engine version to be used by the global database.

Valid values: 1.2.0.0 or above.

Type: String

Required: No

## GlobalClusterIdentifier

The cluster identifier of the new global database cluster.

Type: String

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

Pattern: `[A-Za-z][0-9A-Za-z-:._]*`

Required: Yes

## SourceDBClusterIdentifier

*(Optional)* The Amazon Resource Name (ARN) of an existing Neptune DB cluster to use as the primary cluster of the new global database.

Type: String

Required: No

## StorageEncrypted

The storage encryption setting for the new global database cluster.

Type: Boolean

Required: No

## Response Elements

The following element is returned by the service.

### GlobalCluster

Contains the details of an Amazon Neptune global database.

This data type is used as a response element for the [CreateGlobalCluster](#), [DescribeGlobalClusters](#), [ModifyGlobalCluster](#), [DeleteGlobalCluster](#), [FailoverGlobalCluster](#), and [RemoveFromGlobalCluster](#) actions.

Type: [GlobalCluster](#) object

## Errors

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

### **DBClusterNotFoundFault**

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### **GlobalClusterAlreadyExistsFault**

The `GlobalClusterIdentifier` already exists. Choose a new global database identifier (unique name) to create a new global database cluster.

HTTP Status Code: 400

### **GlobalClusterQuotaExceededFault**

The number of global database clusters for this account is already at the maximum allowed.

HTTP Status Code: 400

### **InvalidDBClusterStateFault**

The DB cluster is not in a valid state.

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 PHP V3](#)

- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# DeleteDBCluster

The DeleteDBCluster action deletes a previously provisioned DB cluster. When you delete a DB cluster, all automated backups for that DB cluster are deleted and can't be recovered. Manual DB cluster snapshots of the specified DB cluster are not deleted.

Note that the DB Cluster cannot be deleted if deletion protection is enabled. To delete it, you must first set its `DeletionProtection` field to `False`.

## Request Parameters

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

### DBClusterIdentifier

The DB cluster identifier for the DB cluster to be deleted. This parameter isn't case-sensitive.

Constraints:

- Must match an existing DBClusterIdentifier.

Type: String

Required: Yes

### FinalDBSnapshotIdentifier

The DB cluster snapshot identifier of the new DB cluster snapshot created when `SkipFinalSnapshot` is set to `false`.

#### Note

Specifying this parameter and also setting the `SkipFinalShapshot` parameter to `true` results in an error.

Constraints:

- Must be 1 to 255 letters, numbers, or hyphens.
- First character must be a letter
- Cannot end with a hyphen or contain two consecutive hyphens

Type: String

Required: No

### **SkipFinalSnapshot**

Determines whether a final DB cluster snapshot is created before the DB cluster is deleted. If `true` is specified, no DB cluster snapshot is created. If `false` is specified, a DB cluster snapshot is created before the DB cluster is deleted.

#### **Note**

You must specify a `FinalDBSnapshotIdentifier` parameter if `SkipFinalSnapshot` is `false`.

Default: `false`

Type: Boolean

Required: No

## **Response Elements**

The following element is returned by the service.

### **DBCluster**

Contains the details of an Amazon Neptune DB cluster.

This data type is used as a response element in the [DescribeDBClusters](#).

Type: [DBCluster](#) object

## **Errors**

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

### **DBClusterNotFoundFault**

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### **DBClusterSnapshotAlreadyExistsFault**

User already has a DB cluster snapshot with the given identifier.

HTTP Status Code: 400

### **InvalidDBClusterSnapshotStateFault**

The supplied value is not a valid DB cluster snapshot state.

HTTP Status Code: 400

### **InvalidDBClusterStateFault**

The DB cluster is not in a valid state.

HTTP Status Code: 400

### **SnapshotQuotaExceeded**

Request would result in user exceeding the allowed number of DB snapshots.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteDBClusterEndpoint

Deletes a custom endpoint and removes it from an Amazon Neptune DB cluster.

## Request Parameters

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

### DBClusterEndpointIdentifier

The identifier associated with the custom endpoint. This parameter is stored as a lowercase string.

Type: String

Required: Yes

## Response Elements

The following elements are returned by the service.

### CustomEndpointType

The type associated with a custom endpoint. One of: READER, WRITER, ANY.

Type: String

### DBClusterEndpointArn

The Amazon Resource Name (ARN) for the endpoint.

Type: String

### DBClusterEndpointIdentifier

The identifier associated with the endpoint. This parameter is stored as a lowercase string.

Type: String

### DBClusterEndpointResourceIdentifier

A unique system-generated identifier for an endpoint. It remains the same for the whole life of the endpoint.

Type: String

### **DBClusterIdentifier**

The DB cluster identifier of the DB cluster associated with the endpoint. This parameter is stored as a lowercase string.

Type: String

### **Endpoint**

The DNS address of the endpoint.

Type: String

### **EndpointType**

The type of the endpoint. One of: `READER`, `WRITER`, `CUSTOM`.

Type: String

### **ExcludedMembers.member.N**

List of DB instance identifiers that aren't part of the custom endpoint group. All other eligible instances are reachable through the custom endpoint. Only relevant if the list of static members is empty.

Type: Array of strings

### **StaticMembers.member.N**

List of DB instance identifiers that are part of the custom endpoint group.

Type: Array of strings

### **Status**

The current status of the endpoint. One of: `creating`, `available`, `deleting`, `inactive`, `modifying`. The `inactive` state applies to an endpoint that cannot be used for a certain kind of cluster, such as a `writer` endpoint for a read-only secondary cluster in a global database.

Type: String

## **Errors**

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

## **DBClusterEndpointNotFoundFault**

The specified custom endpoint doesn't exist.

HTTP Status Code: 400

## **InvalidDBClusterEndpointStateFault**

The requested operation cannot be performed on the endpoint while the endpoint is in this state.

HTTP Status Code: 400

## **InvalidDBClusterStateFault**

The DB cluster is not in a valid state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteDBClusterParameterGroup

Deletes a specified DB cluster parameter group. The DB cluster parameter group to be deleted can't be associated with any DB clusters.

## Request Parameters

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

### DBClusterParameterGroupName

The name of the DB cluster parameter group.

Constraints:

- Must be the name of an existing DB cluster parameter group.
- You can't delete a default DB cluster parameter group.
- Cannot be associated with any DB clusters.

Type: String

Required: Yes

## Errors

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

### DBParameterGroupNotFound

*DBParameterGroupName* does not refer to an existing DB parameter group.

HTTP Status Code: 404

### InvalidDBParameterGroupState

The DB parameter group is in use or is in an invalid state. If you are attempting to delete the parameter group, you cannot delete it when the parameter group is in this state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# DeleteDBClusterSnapshot

Deletes a DB cluster snapshot. If the snapshot is being copied, the copy operation is terminated.

## Note

The DB cluster snapshot must be in the `available` state to be deleted.

## Request Parameters

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

### DBClusterSnapshotIdentifier

The identifier of the DB cluster snapshot to delete.

Constraints: Must be the name of an existing DB cluster snapshot in the `available` state.

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### DBClusterSnapshot

Contains the details for an Amazon Neptune DB cluster snapshot

This data type is used as a response element in the [DescribeDBClusterSnapshots](#) action.

Type: [DBClusterSnapshot](#) object

## Errors

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

## **DBClusterSnapshotNotFoundFault**

*DBClusterSnapshotIdentifier* does not refer to an existing DB cluster snapshot.

HTTP Status Code: 404

## **InvalidDBClusterSnapshotStateFault**

The supplied value is not a valid DB cluster snapshot state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteDBInstance

The DeleteDBInstance action deletes a previously provisioned DB instance. When you delete a DB instance, all automated backups for that instance are deleted and can't be recovered. Manual DB snapshots of the DB instance to be deleted by DeleteDBInstance are not deleted.

If you request a final DB snapshot the status of the Amazon Neptune DB instance is deleting until the DB snapshot is created. The API action DescribeDBInstance is used to monitor the status of this operation. The action can't be canceled or reverted once submitted.

Note that when a DB instance is in a failure state and has a status of `failed`, `incompatible-restore`, or `incompatible-network`, you can only delete it when the `SkipFinalSnapshot` parameter is set to `true`.

You can't delete a DB instance if it is the only instance in the DB cluster, or if it has deletion protection enabled.

## Request Parameters

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

### DBInstanceIdentifier

The DB instance identifier for the DB instance to be deleted. This parameter isn't case-sensitive.

Constraints:

- Must match the name of an existing DB instance.

Type: String

Required: Yes

### FinalDBSnapshotIdentifier

The DBSnapshotIdentifier of the new DBSnapshot created when SkipFinalSnapshot is set to `false`.

#### Note

Specifying this parameter and also setting the SkipFinalShapshot parameter to `true` results in an error.

**Constraints:**

- Must be 1 to 255 letters or numbers.
- First character must be a letter
- Cannot end with a hyphen or contain two consecutive hyphens
- Cannot be specified when deleting a Read Replica.

Type: String


Required: No

**SkipFinalSnapshot**

Determines whether a final DB snapshot is created before the DB instance is deleted. If `true` is specified, no DBSnapshot is created. If `false` is specified, a DB snapshot is created before the DB instance is deleted.

Note that when a DB instance is in a failure state and has a status of 'failed', 'incompatible-restore', or 'incompatible-network', it can only be deleted when the `SkipFinalSnapshot` parameter is set to "true".

Specify `true` when deleting a Read Replica.

 **Note**

The `FinalDBSnapshotIdentifier` parameter must be specified if `SkipFinalSnapshot` is `false`.

Default: `false`

Type: Boolean

Required: No

**Response Elements**

The following element is returned by the service.

**DBInstance**

Contains the details of an Amazon Neptune DB instance.

This data type is used as a response element in the [DescribeDBInstances](#) action.

Type: [DBInstance](#) object

## Errors

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

### **DBInstanceNotFound**

*DBInstanceIdentifier* does not refer to an existing DB instance.

HTTP Status Code: 404

### **DBSnapshotAlreadyExists**

*DBSnapshotIdentifier* is already used by an existing snapshot.

HTTP Status Code: 400

### **InvalidDBClusterStateFault**

The DB cluster is not in a valid state.

HTTP Status Code: 400

### **InvalidDBInstanceState**

The specified DB instance is not in the *available* state.

HTTP Status Code: 400

### **SnapshotQuotaExceeded**

Request would result in user exceeding the allowed number of DB snapshots.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteDBParameterGroup

Deletes a specified DBParameterGroup. The DBParameterGroup to be deleted can't be associated with any DB instances.

## Request Parameters

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

### DBParameterGroupName

The name of the DB parameter group.

Constraints:

- Must be the name of an existing DB parameter group
- You can't delete a default DB parameter group
- Cannot be associated with any DB instances

Type: String

Required: Yes

## Errors

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

### DBParameterGroupNotFound

*DBParameterGroupName* does not refer to an existing DB parameter group.

HTTP Status Code: 404

### InvalidDBParameterGroupState

The DB parameter group is in use or is in an invalid state. If you are attempting to delete the parameter group, you cannot delete it when the parameter group is in this state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# DeleteDBSubnetGroup

Deletes a DB subnet group.

## Note

The specified database subnet group must not be associated with any DB instances.

## Request Parameters

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

### DBSubnetGroupName

The name of the database subnet group to delete.

## Note

You can't delete the default subnet group.

Constraints:

Constraints: Must match the name of an existing DBSubnetGroup. Must not be default.

Example: mySubnetgroup

Type: String

Required: Yes

## Errors

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

### DBSubnetGroupNotFoundFault

*DBSubnetGroupName* does not refer to an existing DB subnet group.

HTTP Status Code: 404

### **InvalidDBSubnetGroupStateFault**

The DB subnet group cannot be deleted because it is in use.

HTTP Status Code: 400

### **InvalidDBSubnetStateFault**

The DB subnet is not in the *available* state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteEventSubscription

Deletes an event notification subscription.

## Request Parameters

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

### SubscriptionName

The name of the event notification subscription you want to delete.

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### EventSubscription

Contains the results of a successful invocation of the [DescribeEventSubscriptions](#) action.

Type: [EventSubscription](#) object

## Errors

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

### InvalidEventSubscriptionState

The event subscription is in an invalid state.

HTTP Status Code: 400

### SubscriptionNotFound

The designated subscription 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteGlobalCluster

Deletes a global database. The primary and all secondary clusters must already be detached or deleted first.

## Request Parameters

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

### GlobalClusterIdentifier

The cluster identifier of the global database cluster being deleted.

Type: String

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

Pattern: `[A-Za-z][0-9A-Za-z-:._]*`

Required: Yes

## Response Elements

The following element is returned by the service.

### GlobalCluster

Contains the details of an Amazon Neptune global database.

This data type is used as a response element for the [CreateGlobalCluster](#), [DescribeGlobalClusters](#), [ModifyGlobalCluster](#), [DeleteGlobalCluster](#), [FailoverGlobalCluster](#), and [RemoveFromGlobalCluster](#) actions.

Type: [GlobalCluster](#) object

## Errors

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

### GlobalClusterNotFoundFault

The `GlobalClusterIdentifier` doesn't refer to an existing global database cluster.

HTTP Status Code: 404

## **InvalidGlobalClusterStateFault**

The global cluster is in an invalid state and can't perform the requested operation.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeDBClusterEndpoints

Returns information about endpoints for an Amazon Neptune DB cluster.

## Note

This operation can also return information for Amazon RDS clusters and Amazon DocDB clusters.

## Request Parameters

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

### DBClusterEndpointIdentifier

The identifier of the endpoint to describe. This parameter is stored as a lowercase string.

Type: String

Required: No

### DBClusterIdentifier

The DB cluster identifier of the DB cluster associated with the endpoint. This parameter is stored as a lowercase string.

Type: String

Required: No

### Filters.Filter.N

A set of name-value pairs that define which endpoints to include in the output. The filters are specified as name-value pairs, in the format `Name=endpoint_type, Values=endpoint_type1, endpoint_type2, . . .`. Name can be one of: `db-cluster-endpoint-type`, `db-cluster-endpoint-custom-type`, `db-cluster-endpoint-id`, `db-cluster-endpoint-status`. Values for the `db-cluster-endpoint-type` filter can be one or more of: `reader`, `writer`, `custom`. Values for the `db-cluster-endpoint-custom-type` filter can be one or more of: `reader`, `any`. Values for the `db-cluster-endpoint-status` filter can be one or more of: `available`, `creating`, `deleting`, `inactive`, `modifying`.

Type: Array of [Filter](#) objects

Required: No

### Marker

An optional pagination token provided by a previous `DescribeDBClusterEndpoints` request. 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 pagination token called a marker is included in the response so you can retrieve the remaining results.

Default: 100

Constraints: Minimum 20, maximum 100.

Type: Integer

Required: No

## Response Elements

The following elements are returned by the service.

### **DBClusterEndpoints.DBClusterEndpointList.N**

Contains the details of the endpoints associated with the cluster and matching any filter conditions.

Type: Array of [DBClusterEndpoint](#) objects

### Marker

An optional pagination token provided by a previous `DescribeDBClusterEndpoints` request. 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](#).

### **DBClusterNotFoundFault**

*DBClusterIdentifier* does not refer to an existing DB cluster.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeDBClusterParameterGroups

Returns a list of `DBClusterParameterGroup` descriptions. If a `DBClusterParameterGroupName` parameter is specified, the list will contain only the description of the specified DB cluster parameter group.

## Request Parameters

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

### `DBClusterParameterGroupName`

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

Constraints:

- If supplied, must match the name of an existing `DBClusterParameterGroup`.

Type: String

Required: No

### `Filters.Filter.N`

This parameter is not currently supported.

Type: Array of [Filter](#) objects

Required: No

### `Marker`

An optional pagination token provided by a previous `DescribeDBClusterParameterGroups` request. 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 pagination token called 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.

### **DBClusterParameterGroups.DBClusterParameterGroup.N**

A list of DB cluster parameter groups.

Type: Array of [DBClusterParameterGroup](#) objects

### **Marker**

An optional pagination token provided by a previous `DescribeDBClusterParameterGroups` request. 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](#).

### **DBParameterGroupNotFound**

*DBParameterGroupName* does not refer to an existing DB parameter group.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeDBClusterParameters

Returns the detailed parameter list for a particular DB cluster parameter group.

## Request Parameters

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

### DBClusterParameterGroupName

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

Constraints:

- If supplied, must match the name of an existing DBClusterParameterGroup.

Type: String

Required: Yes

### Filters.Filter.N

This parameter is not currently supported.

Type: Array of [Filter](#) objects

Required: No

### Marker

An optional pagination token provided by a previous DescribeDBClusterParameters request. 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 pagination token called 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

A value that indicates to return only parameters for a specific source. Parameter sources can be engine, service, or customer.

Type: String

Required: No

## Response Elements

The following elements are returned by the service.

### Marker

An optional pagination token provided by a previous DescribeDBClusterParameters request. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by MaxRecords .

Type: String

### Parameters.Parameter.N

Provides a list of parameters for the DB cluster parameter group.

Type: Array of [Parameter](#) objects

## Errors

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

### DBParameterGroupNotFound

*DBParameterGroupName* does not refer to an existing DB parameter group.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeDBClusters

Returns information about provisioned DB clusters, and supports pagination.

## Note

This operation can also return information for Amazon RDS clusters and Amazon DocDB clusters.

## Request Parameters

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

### DBClusterIdentifier

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

Constraints:

- If supplied, must match an existing DBClusterIdentifier.

Type: String

Required: No

### Filters.Filter.N

A filter that specifies one or more DB clusters to describe.

Supported filters:

- `db-cluster-id` - Accepts DB cluster identifiers and DB cluster Amazon Resource Names (ARNs). The results list will only include information about the DB clusters identified by these ARNs.
- `engine` - Accepts an engine name (such as `neptune`), and restricts the results list to DB clusters created by that engine.

For example, to invoke this API from the Amazon CLI and filter so that only Neptune DB clusters are returned, you could use the following command:



Type: Array of [Filter](#) objects

Required: No

### Marker

An optional pagination token provided by a previous [DescribeDBClusters](#) request. 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 pagination token called 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.

### DBClusters.DBCluster.N

Contains a list of DB clusters for the user.

Type: Array of [DBCluster](#) objects

### Marker

A pagination token that can be used in a subsequent DescribeDBClusters request.

Type: String

## Errors

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

### **DBClusterNotFoundFault**

*DBClusterIdentifier* does not refer to an existing DB cluster.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeDBClusterSnapshotAttributes

Returns a list of DB cluster snapshot attribute names and values for a manual DB cluster snapshot.

When sharing snapshots with other Amazon accounts, `DescribeDBClusterSnapshotAttributes` returns the `restore` attribute and a list of IDs for the Amazon accounts that are authorized to copy or restore the manual DB cluster snapshot. If `all` is included in the list of values for the `restore` attribute, then the manual DB cluster snapshot is public and can be copied or restored by all Amazon accounts.

To add or remove access for an Amazon account to copy or restore a manual DB cluster snapshot, or to make the manual DB cluster snapshot public or private, use the [ModifyDBClusterSnapshotAttribute](#) API action.

## Request Parameters

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

### DBClusterSnapshotIdentifier

The identifier for the DB cluster snapshot to describe the attributes for.

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### DBClusterSnapshotAttributesResult

Contains the results of a successful call to the [DescribeDBClusterSnapshotAttributes](#) API action.

Manual DB cluster snapshot attributes are used to authorize other Amazon accounts to copy or restore a manual DB cluster snapshot. For more information, see the [ModifyDBClusterSnapshotAttribute](#) API action.

Type: [DBClusterSnapshotAttributesResult](#) object

## Errors

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

### **DBClusterSnapshotNotFoundFault**

*DBClusterSnapshotIdentifier* does not refer to an existing DB cluster snapshot.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeDBClusterSnapshots

Returns information about DB cluster snapshots. This API action supports pagination.

## Request Parameters

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

### DBClusterIdentifier

The ID of the DB cluster to retrieve the list of DB cluster snapshots for. This parameter can't be used in conjunction with the `DBClusterSnapshotIdentifier` parameter. This parameter is not case-sensitive.

Constraints:

- If supplied, must match the identifier of an existing `DBCluster`.

Type: String

Required: No

### DBClusterSnapshotIdentifier

A specific DB cluster snapshot identifier to describe. This parameter can't be used in conjunction with the `DBClusterIdentifier` parameter. This value is stored as a lowercase string.

Constraints:

- If supplied, must match the identifier of an existing `DBClusterSnapshot`.
- If this identifier is for an automated snapshot, the `SnapshotType` parameter must also be specified.

Type: String

Required: No

### Filters.Filter.N

This parameter is not currently supported.

Type: Array of [Filter](#) objects

Required: No

### **IncludePublic**

True to include manual DB cluster snapshots that are public and can be copied or restored by any Amazon account, and otherwise false. The default is `false`. The default is false.

You can share a manual DB cluster snapshot as public by using the [ModifyDBClusterSnapshotAttribute](#) API action.

Type: Boolean

Required: No

### **IncludeShared**

True to include shared manual DB cluster snapshots from other Amazon accounts that this Amazon account has been given permission to copy or restore, and otherwise false. The default is `false`.

You can give an Amazon account permission to restore a manual DB cluster snapshot from another Amazon account by the [ModifyDBClusterSnapshotAttribute](#) API action.

Type: Boolean

Required: No

### **Marker**

An optional pagination token provided by a previous `DescribeDBClusterSnapshots` request. 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 pagination token called 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

### SnapshotType

The type of DB cluster snapshots to be returned. You can specify one of the following values:

- `automated` - Return all DB cluster snapshots that have been automatically taken by Amazon Neptune for my Amazon account.
- `manual` - Return all DB cluster snapshots that have been taken by my Amazon account.
- `shared` - Return all manual DB cluster snapshots that have been shared to my Amazon account.
- `public` - Return all DB cluster snapshots that have been marked as public.

If you don't specify a `SnapshotType` value, then both `automated` and `manual` DB cluster snapshots are returned. You can include shared DB cluster snapshots with these results by setting the `IncludeShared` parameter to `true`. You can include public DB cluster snapshots with these results by setting the `IncludePublic` parameter to `true`.

The `IncludeShared` and `IncludePublic` parameters don't apply for `SnapshotType` values of `manual` or `automated`. The `IncludePublic` parameter doesn't apply when `SnapshotType` is set to `shared`. The `IncludeShared` parameter doesn't apply when `SnapshotType` is set to `public`.

Type: String

Required: No

## Response Elements

The following elements are returned by the service.

### **DBClusterSnapshots.DBClusterSnapshot.N**

Provides a list of DB cluster snapshots for the user.

Type: Array of [DBClusterSnapshot](#) objects

## Marker

An optional pagination token provided by a previous [DescribeDBClusterSnapshots](#) request. 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](#).

### **DBClusterSnapshotNotFoundFault**

*DBClusterSnapshotIdentifier* does not refer to an existing DB cluster snapshot.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# DescribeDBEngineVersions

Returns a list of the available DB engines.

## Request Parameters

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

### DBParameterGroupFamily

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

Constraints:

- If supplied, must match an existing DBParameterGroupFamily.

Type: String

Required: No

### DefaultOnly

Indicates that only the default version of the specified engine or engine and major version combination is returned.

Type: Boolean

Required: No

### Engine

The database engine to return.

Type: String

Required: No

### EngineVersion

The database engine version to return.

Example: 5.1.49

Type: String

Required: No

## Filters.Filter.N

Not currently supported.

Type: Array of [Filter](#) objects

Required: No

## ListSupportedCharacterSets

If this parameter is specified and the requested engine supports the `CharacterSetName` parameter for `CreateDBInstance`, the response includes a list of supported character sets for each engine version.

Type: Boolean

Required: No

## ListSupportedTimezones

If this parameter is specified and the requested engine supports the `TimeZone` parameter for `CreateDBInstance`, the response includes a list of supported time zones for each engine version.

Type: Boolean

Required: No

## Marker

An optional pagination token provided by a previous request. 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 than the `MaxRecords` value is available, a pagination token called a marker is included in the response so that the following 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.

### **DBEngineVersions.DBEngineVersion.N**

A list of `DBEngineVersion` elements.

Type: Array of [DBEngineVersion](#) objects

### **Marker**

An optional pagination token provided by a previous request. 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](#).

## 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 PHP V3](#)

- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeDBInstances

Returns information about provisioned instances, and supports pagination.

## Note

This operation can also return information for Amazon RDS instances and Amazon DocDB instances.

## Request Parameters

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

### DBInstanceIdentifier

The user-supplied instance identifier. If this parameter is specified, information from only the specific DB instance is returned. This parameter isn't case-sensitive.

Constraints:

- If supplied, must match the identifier of an existing DBInstance.

Type: String

Required: No

### Filters.Filter.N

A filter that specifies one or more DB instances to describe.

Supported filters:

- `db-cluster-id` - Accepts DB cluster identifiers and DB cluster Amazon Resource Names (ARNs). The results list will only include information about the DB instances associated with the DB clusters identified by these ARNs.
- `engine` - Accepts an engine name (such as `neptune`), and restricts the results list to DB instances created by that engine.

For example, to invoke this API from the Amazon CLI and filter so that only Neptune DB instances are returned, you could use the following command:

Type: Array of [Filter](#) objects

Required: No

### Marker

An optional pagination token provided by a previous `DescribeDBInstances` request. 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 pagination token called 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.

### DBInstances.DBInstance.N

A list of [DBInstance](#) instances.

Type: Array of [DBInstance](#) objects

### Marker

An optional pagination token provided by a previous request. 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](#).

### **DBInstanceNotFound**

*DBInstanceIdentifier* does not refer to an existing DB instance.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeDBParameterGroups

Returns a list of `DBParameterGroup` descriptions. If a `DBParameterGroupName` is specified, the list will contain only the description of the specified DB parameter group.

## Request Parameters

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

### **DBParameterGroupName**

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

Constraints:

- If supplied, must match the name of an existing `DBClusterParameterGroup`.

Type: String

Required: No

### **Filters.Filter.N**

This parameter is not currently supported.

Type: Array of [Filter](#) objects

Required: No

### **Marker**

An optional pagination token provided by a previous `DescribeDBParameterGroups` request. 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 pagination token called 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.

### **DBParameterGroups.DBParameterGroup.N**

A list of [DBParameterGroup](#) instances.

Type: Array of [DBParameterGroup](#) objects

### **Marker**

An optional pagination token provided by a previous request. 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](#).

### **DBParameterGroupNotFound**

*DBParameterGroupName* does not refer to an existing DB parameter group.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeDBParameters

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

## Request Parameters

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

### DBParameterGroupName

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

Constraints:

- If supplied, must match the name of an existing DBParameterGroup.

Type: String

Required: Yes

### Filters.Filter.N

This parameter is not currently supported.

Type: Array of [Filter](#) objects

Required: No

### Marker

An optional pagination token provided by a previous DescribeDBParameters request. 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 pagination token called 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.

Default: All parameter types returned

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

Type: String

Required: No

## Response Elements

The following elements are returned by the service.

### Marker

An optional pagination token provided by a previous request. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

Type: String

### Parameters.Parameter.N

A list of [Parameter](#) values.

Type: Array of [Parameter](#) objects

## Errors

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

### DBParameterGroupNotFound

*DBParameterGroupName* does not refer to an existing DB parameter group.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeDBSubnetGroups

Returns a list of DBSubnetGroup descriptions. If a DBSubnetGroupName is specified, the list will contain only the descriptions of the specified DBSubnetGroup.

For an overview of CIDR ranges, go to the [Wikipedia Tutorial](#).

## Request Parameters

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

### DBSubnetGroupName

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

Type: String

Required: No

### Filters.Filter.N

This parameter is not currently supported.

Type: Array of [Filter](#) objects

Required: No

### Marker

An optional pagination token provided by a previous DescribeDBSubnetGroups request. 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 pagination token called 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.

### **DBSubnetGroups.DBSubnetGroup.N**

A list of [DBSubnetGroup](#) instances.

Type: Array of [DBSubnetGroup](#) objects

### **Marker**

An optional pagination token provided by a previous request. 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](#).

### **DBSubnetGroupNotFoundFault**

*DBSubnetGroupName* does not refer to an existing DB subnet group.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# DescribeEngineDefaultClusterParameters

Returns the default engine and system parameter information for the cluster database engine.

## Request Parameters

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

### DBParameterGroupFamily

The name of the DB cluster parameter group family to return engine parameter information for.

Type: String

Required: Yes

### Filters.Filter.N

This parameter is not currently supported.

Type: Array of [Filter](#) objects

Required: No

### Marker

An optional pagination token provided by a previous `DescribeEngineDefaultClusterParameters` request. 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 pagination token called 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

Contains the result of a successful invocation of the [DescribeEngineDefaultParameters](#) action.

Type: [EngineDefaults](#) object

## Errors

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

## 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeEngineDefaultParameters

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

## Request Parameters

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

### DBParameterGroupFamily

The name of the DB parameter group family.

Type: String

Required: Yes

### Filters.Filter.N

Not currently supported.

Type: Array of [Filter](#) objects

Required: No

### Marker

An optional pagination token provided by a previous DescribeEngineDefaultParameters request. 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 pagination token called 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

Contains the result of a successful invocation of the [DescribeEngineDefaultParameters](#) action.

Type: [EngineDefaults](#) object

## Errors

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

## 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeEventCategories

Displays a list of categories for all event source types, or, if specified, for a specified source type.

## Request Parameters

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

### Filters.Filter.N

This parameter is not currently supported.

Type: Array of [Filter](#) objects

Required: No

### SourceType

The type of source that is generating the events.

Valid values: db-instance | db-parameter-group | db-security-group | db-snapshot

Type: String

Required: No

## Response Elements

The following element is returned by the service.

### EventCategoriesMapList.EventCategoriesMap.N

A list of EventCategoriesMap data types.

Type: Array of [EventCategoriesMap](#) objects

## Errors

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

## 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeEvents

Returns events related to DB instances, DB security groups, DB snapshots, and DB parameter groups for the past 14 days. Events specific to a particular DB instance, DB security group, database snapshot, or DB parameter group can be obtained by providing the name as a parameter. By default, the past hour of events are returned.

## Request Parameters

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

### Duration

The number of minutes to retrieve events for.

Default: 60

Type: Integer

Required: No

### EndTime

The end of the time interval for which to retrieve events, specified in ISO 8601 format. For more information about ISO 8601, go to the [ISO8601 Wikipedia page](#).

Example: 2009-07-08T18:00Z

Type: Timestamp

Required: No

### EventCategories.EventCategory.N

A list of event categories that trigger notifications for a event notification subscription.

Type: Array of strings

Required: No

### Filters.Filter.N

This parameter is not currently supported.

Type: Array of [Filter](#) objects

Required: No

### Marker

An optional pagination token provided by a previous DescribeEvents request. 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 pagination token called 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, then all sources are included in the response.

Constraints:

- If SourceIdentifier is supplied, SourceType must also be provided.
- If the source type is DBInstance, then a DBInstanceIdentifier must be supplied.
- If the source type is DBSecurityGroup, a DBSecurityGroupName must be supplied.
- If the source type is DBParameterGroup, a DBParameterGroupName must be supplied.
- If the source type is DBSnapshot, a DBSnapshotIdentifier must be supplied.
- Cannot end with a hyphen or contain two consecutive hyphens.

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: db-instance | db-parameter-group | db-security-group | db-snapshot | db-cluster | db-cluster-snapshot

Required: No

## StartTime

The beginning of the time interval to retrieve events for, specified in ISO 8601 format. For more information about ISO 8601, go to the [ISO8601 Wikipedia page](#).

Example: 2009-07-08T18:00Z

Type: Timestamp

Required: No

## Response Elements

The following elements are returned by the service.

### Events.Event.N

A list of [Event](#) instances.

Type: Array of [Event](#) objects

### Marker

An optional pagination token provided by a previous Events request. 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](#).

## 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeEventSubscriptions

Lists all the subscription descriptions for a customer account. The description for a subscription includes SubscriptionName, SNSTopicARN, CustomerID, SourceType, SourceID, CreationTime, and Status.

If you specify a SubscriptionName, lists the description for that subscription.

## Request Parameters

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

### Filters.Filter.N

This parameter is not currently supported.

Type: Array of [Filter](#) objects

Required: No

### Marker

An optional pagination token provided by a previous DescribeOrderableDBInstanceOptions request. 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 pagination token called 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

## SubscriptionName

The name of the event notification subscription you want to describe.

Type: String

Required: No

## Response Elements

The following elements are returned by the service.

### EventSubscriptionsList.EventSubscription.N

A list of EventSubscriptions data types.

Type: Array of [EventSubscription](#) objects

### Marker

An optional pagination token provided by a previous DescribeOrderableDBInstanceOptions request. 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](#).

### SubscriptionNotFound

The designated subscription 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeGlobalClusters

Returns information about Neptune global database clusters. This API supports pagination.

## Request Parameters

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

### GlobalClusterIdentifier

The user-supplied DB cluster identifier. If this parameter is specified, only information about the specified DB cluster is returned. This parameter is not case-sensitive.

Constraints: If supplied, must match an existing DB cluster identifier.

Type: String

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

Pattern: `[A-Za-z][0-9A-Za-z-:._]*`

Required: No

### Marker

*(Optional)* A pagination token returned by a previous call to `DescribeGlobalClusters`. If this parameter is specified, the response will only include records beyond the marker, up to the number 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 pagination marker token is included in the response that you can use to retrieve the remaining results.

Default: 100

Constraints: Minimum 20, maximum 100.

Type: Integer

Required: No

## Response Elements

The following elements are returned by the service.

### **GlobalClusters.GlobalClusterMember.N**

The list of global clusters and instances returned by this request.

Type: Array of [GlobalCluster](#) objects

### **Marker**

A pagination token. If this parameter is returned in the response, more records are available, which can be retrieved by one or more additional calls to `DescribeGlobalClusters`.

Type: String

## Errors

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

### **GlobalClusterNotFoundFault**

The `GlobalClusterIdentifier` doesn't refer to an existing global database cluster.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# DescribeOrderableDBInstanceOptions

Returns a list of orderable DB instance options for the specified engine.

## Request Parameters

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

### DBInstanceClass

The DB instance class filter value. Specify this parameter to show only the available offerings matching the specified DB instance class.

Type: String

Required: No

### Engine

The name of the engine to retrieve DB instance options for.

Type: String

Required: Yes

### EngineVersion

The engine version filter value. Specify this parameter to show only the available offerings matching the specified engine version.

Type: String

Required: No

### Filters.Filter.N

This parameter is not currently supported.

Type: Array of [Filter](#) objects

Required: No

### LicenseModel

The license model filter value. Specify this parameter to show only the available offerings matching the specified license model.

Type: String

Required: No

### Marker

An optional pagination token provided by a previous `DescribeOrderableDBInstanceOptions` request. 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 pagination token called 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

### Vpc

The VPC filter value. Specify this parameter to show only the available VPC or non-VPC offerings.

Type: Boolean

Required: No

## Response Elements

The following elements are returned by the service.

## Marker

An optional pagination token provided by a previous `OrderableDBInstanceOptions` request. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords` .

Type: String

## `OrderableDBInstanceOptions.OrderableDBInstanceOption.N`

An [OrderableDBInstanceOption](#) structure containing information about orderable options for the DB instance.

Type: Array of [OrderableDBInstanceOption](#) objects

## Errors

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

## 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribePendingMaintenanceActions

Returns a list of resources (for example, DB instances) that have at least one pending maintenance action.

## Request Parameters

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

### Filters.Filter.N

A filter that specifies one or more resources to return pending maintenance actions for.

Supported filters:

- `db-cluster-id` - Accepts DB cluster identifiers and DB cluster Amazon Resource Names (ARNs). The results list will only include pending maintenance actions for the DB clusters identified by these ARNs.
- `db-instance-id` - Accepts DB instance identifiers and DB instance ARNs. The results list will only include pending maintenance actions for the DB instances identified by these ARNs.

Type: Array of [Filter](#) objects

Required: No

### Marker

An optional pagination token provided by a previous `DescribePendingMaintenanceActions` request. If this parameter is specified, the response includes only records beyond the marker, up to a number of records 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 pagination token called 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

### **ResourceIdentifier**

The ARN of a resource to return pending maintenance actions for.

Type: String

Required: No

## **Response Elements**

The following elements are returned by the service.

### **Marker**

An optional pagination token provided by a previous `DescribePendingMaintenanceActions` request. If this parameter is specified, the response includes only records beyond the marker, up to a number of records specified by `MaxRecords`.

Type: String

### **PendingMaintenanceActions.ResourcePendingMaintenanceActions.N**

A list of the pending maintenance actions for the resource.

Type: Array of [ResourcePendingMaintenanceActions](#) objects

## **Errors**

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

### **ResourceNotFoundFault**

The specified resource ID was not 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeValidDBInstanceModifications

You can call [DescribeValidDBInstanceModifications](#) to learn what modifications you can make to your DB instance. You can use this information when you call [ModifyDBInstance](#).

## Request Parameters

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

### DBInstanceIdentifier

The customer identifier or the ARN of your DB instance.

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### ValidDBInstanceModificationsMessage

Information about valid modifications that you can make to your DB instance. Contains the result of a successful call to the [DescribeValidDBInstanceModifications](#) action. You can use this information when you call [ModifyDBInstance](#).

Type: [ValidDBInstanceModificationsMessage](#) object

## Errors

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

### DBInstanceNotFound

*DBInstanceIdentifier* does not refer to an existing DB instance.

HTTP Status Code: 404

### InvalidDBInstanceState

The specified DB instance is not in the *available* state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# FailoverDBCluster

Forces a failover for a DB cluster.

A failover for a DB cluster promotes one of the Read Replicas (read-only instances) in the DB cluster to be the primary instance (the cluster writer).

Amazon Neptune will automatically fail over to a Read Replica, if one exists, when the primary instance fails. You can force a failover when you want to simulate a failure of a primary instance for testing. Because each instance in a DB cluster has its own endpoint address, you will need to clean up and re-establish any existing connections that use those endpoint addresses when the failover is complete.

## Request Parameters

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

### DBClusterIdentifier

A DB cluster identifier to force a failover for. This parameter is not case-sensitive.

Constraints:

- Must match the identifier of an existing DBCluster.

Type: String

Required: No

### TargetDBInstanceIdentifier

The name of the instance to promote to the primary instance.

You must specify the instance identifier for an Read Replica in the DB cluster. For example, `mydbcluster-replica1`.

Type: String

Required: No

## Response Elements

The following element is returned by the service.

## DBCluster

Contains the details of an Amazon Neptune DB cluster.

This data type is used as a response element in the [DescribeDBClusters](#).

Type: [DBCluster](#) object

## Errors

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

### DBClusterNotFoundFault

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### InvalidDBClusterStateFault

The DB cluster is not in a valid state.

HTTP Status Code: 400

### InvalidDBInstanceState

The specified DB instance is not in the *available* state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# FailoverGlobalCluster

Initiates the failover process for a Neptune global database.

A failover for a Neptune global database promotes one of secondary read-only DB clusters to be the primary DB cluster and demotes the primary DB cluster to being a secondary (read-only) DB cluster. In other words, the role of the current primary DB cluster and the selected target secondary DB cluster are switched. The selected secondary DB cluster assumes full read/write capabilities for the Neptune global database.

## Note

This action applies **only** to Neptune global databases. This action is only intended for use on healthy Neptune global databases with healthy Neptune DB clusters and no region-wide outages, to test disaster recovery scenarios or to reconfigure the global database topology.

## Request Parameters

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

### GlobalClusterIdentifier

Identifier of the Neptune global database that should be failed over. The identifier is the unique key assigned by the user when the Neptune global database was created. In other words, it's the name of the global database that you want to fail over.

Constraints: Must match the identifier of an existing Neptune global database.

Type: String

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

Pattern: `[A-Za-z][0-9A-Za-z- : . _]*`

Required: Yes

### TargetDbClusterIdentifier

The Amazon Resource Name (ARN) of the secondary Neptune DB cluster that you want to promote to primary for the global database.

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### GlobalCluster

Contains the details of an Amazon Neptune global database.

This data type is used as a response element for the [CreateGlobalCluster](#), [DescribeGlobalClusters](#), [ModifyGlobalCluster](#), [DeleteGlobalCluster](#), [FailoverGlobalCluster](#), and [RemoveFromGlobalCluster](#) actions.

Type: [GlobalCluster](#) object

## Errors

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

### DBClusterNotFoundFault

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### GlobalClusterNotFoundFault

The `GlobalClusterIdentifier` doesn't refer to an existing global database cluster.

HTTP Status Code: 404

### InvalidDBClusterStateFault

The DB cluster is not in a valid state.

HTTP Status Code: 400

### InvalidGlobalClusterStateFault

The global cluster is in an invalid state and can't perform the requested operation.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ListTagsForResource

Lists all tags on an Amazon Neptune resource.

## Request Parameters

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

### Filters.Filter.N

This parameter is not currently supported.

Type: Array of [Filter](#) objects

Required: No

### ResourceName

The Amazon Neptune resource with tags to be listed. This value is an Amazon Resource Name (ARN). For information about creating an ARN, see [Constructing an Amazon Resource Name \(ARN\)](#).

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### TagList.Tag.N

List of tags returned by the ListTagsForResource operation.

Type: Array of [Tag](#) objects

## Errors

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

## **DBClusterNotFoundFault**

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

## **DBInstanceNotFound**

*DBInstanceIdentifier* does not refer to an existing DB instance.

HTTP Status Code: 404

## **DBSnapshotNotFound**

*DBSnapshotIdentifier* does not refer to an existing DB snapshot.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# ModifyDBCluster

Modify a setting for a DB cluster. You can change one or more database configuration parameters by specifying these parameters and the new values in the request.

## Request Parameters

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

### AllowMajorVersionUpgrade

A value that indicates whether upgrades between different major versions are allowed.

**Constraints:** You must set the allow-major-version-upgrade flag when providing an `EngineVersion` parameter that uses a different major version than the DB cluster's current version.

Type: Boolean

Required: No

### ApplyImmediately

A value that specifies whether the modifications in this request and any pending modifications are asynchronously applied as soon as possible, regardless of the `PreferredMaintenanceWindow` setting for the DB cluster. If this parameter is set to false, changes to the DB cluster are applied during the next maintenance window.

The `ApplyImmediately` parameter only affects `NewDBClusterIdentifier` values. If you set the `ApplyImmediately` parameter value to false, then changes to `NewDBClusterIdentifier` values are applied during the next maintenance window. All other changes are applied immediately, regardless of the value of the `ApplyImmediately` parameter.

Default: false

Type: Boolean

Required: No

### BackupRetentionPeriod

The number of days for which automated backups are retained. You must specify a minimum value of 1.

Default: 1

Constraints:

- Must be a value from 1 to 35

Type: Integer

Required: No

### **CloudwatchLogsExportConfiguration**

The configuration setting for the log types to be enabled for export to CloudWatch Logs for a specific DB cluster. See [Using the CLI to publish Neptune audit logs to CloudWatch Logs](#).

Type: [CloudwatchLogsExportConfiguration](#) object

Required: No

### **CopyTagsToSnapshot**

*If set to `true`, tags are copied to any snapshot of the DB cluster that is created.*

Type: Boolean

Required: No

### **DBClusterIdentifier**

The DB cluster identifier for the cluster being modified. This parameter is not case-sensitive.

Constraints:

- Must match the identifier of an existing DBCluster.

Type: String

Required: Yes

### **DBClusterParameterGroupName**

The name of the DB cluster parameter group to use for the DB cluster.

Type: String

Required: No

## DBInstanceParameterGroupName

The name of the DB parameter group to apply to all instances of the DB cluster.

### Note

When you apply a parameter group using `DBInstanceParameterGroupName`, parameter changes aren't applied during the next maintenance window but instead are applied immediately.

Default: The existing name setting

Constraints:

- The DB parameter group must be in the same DB parameter group family as the target DB cluster version.
- The `DBInstanceParameterGroupName` parameter is only valid in combination with the `AllowMajorVersionUpgrade` parameter.

Type: String

Required: No

## DeletionProtection

A value that indicates whether the DB cluster has deletion protection enabled. The database can't be deleted when deletion protection is enabled. By default, deletion protection is disabled.

Type: Boolean

Required: No

## EnableIAMDatabaseAuthentication

True to enable mapping of Amazon Identity and Access Management (IAM) accounts to database accounts, and otherwise false.

Default: false

Type: Boolean

Required: No

## EngineVersion

The version number of the database engine to which you want to upgrade. Changing this parameter results in an outage. The change is applied during the next maintenance window unless the `ApplyImmediately` parameter is set to true.

For a list of valid engine versions, see [Engine Releases for Amazon Neptune](#), or call [DescribeDBEngineVersions](#).

Type: String

Required: No

## MasterUserPassword

Not supported by Neptune.

Type: String

Required: No

## NewDBClusterIdentifier

The new DB cluster identifier for the DB cluster when renaming a DB cluster. This value is stored as a lowercase string.

Constraints:

- Must contain from 1 to 63 letters, numbers, or hyphens
- The first character must be a letter
- Cannot end with a hyphen or contain two consecutive hyphens

Example: `my-cluster2`

Type: String

Required: No

## OptionGroupName

*Not supported by Neptune.*

Type: String

Required: No

## Port

The port number on which the DB cluster accepts connections.

Constraints: Value must be 1150-65535

Default: The same port as the original DB cluster.

Type: Integer

Required: No

## PreferredBackupWindow

The daily time range during which automated backups are created if automated backups are enabled, using the `BackupRetentionPeriod` parameter.

The default is a 30-minute window selected at random from an 8-hour block of time for each Amazon Region.

Constraints:

- Must be in the format `hh24:mi-hh24:mi`.
- Must be in Universal Coordinated Time (UTC).
- Must not conflict with the preferred maintenance window.
- Must be at least 30 minutes.

Type: String

Required: No

## PreferredMaintenanceWindow

The weekly time range during which system maintenance can occur, in Universal Coordinated Time (UTC).

Format: `ddd:hh24:mi-ddd:hh24:mi`

The default is a 30-minute window selected at random from an 8-hour block of time for each Amazon Region, occurring on a random day of the week.

Valid Days: Mon, Tue, Wed, Thu, Fri, Sat, Sun.

Constraints: Minimum 30-minute window.

Type: String

Required: No

### ServerlessV2ScalingConfiguration

Contains the scaling configuration of a Neptune Serverless DB cluster.

For more information, see [Using Amazon Neptune Serverless](#) in the *Amazon Neptune User Guide*.

Type: [ServerlessV2ScalingConfiguration](#) object

Required: No

### StorageType

The storage type to associate with the DB cluster.

Valid Values:

- **standard** – ( *the default* ) Configures cost-effective database storage for applications with moderate to small I/O usage.
- **iopt1** – Enables [I/O-Optimized storage](#) that's designed to meet the needs of I/O-intensive graph workloads that require predictable pricing with low I/O latency and consistent I/O throughput.

Neptune I/O-Optimized storage is only available starting with engine release 1.3.0.0.

Type: String

Required: No

### VpcSecurityGroupIds.VpcSecurityGroupId.N

A list of VPC security groups that the DB cluster will belong to.

Type: Array of strings

Required: No

## Response Elements

The following element is returned by the service.

## DBCluster

Contains the details of an Amazon Neptune DB cluster.

This data type is used as a response element in the [DescribeDBClusters](#).

Type: [DBCluster](#) object

## Errors

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

### DBClusterAlreadyExistsFault

User already has a DB cluster with the given identifier.

HTTP Status Code: 400

### DBClusterNotFoundFault

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### DBClusterParameterGroupNotFound

*DBClusterParameterGroupName* does not refer to an existing DB Cluster parameter group.

HTTP Status Code: 404

### DBSubnetGroupNotFoundFault

*DBSubnetGroupName* does not refer to an existing DB subnet group.

HTTP Status Code: 404

### InvalidDBClusterStateFault

The DB cluster is not in a valid state.

HTTP Status Code: 400

### InvalidDBInstanceState

The specified DB instance is not in the *available* state.

HTTP Status Code: 400

### **InvalidDBSecurityGroupState**

The state of the DB security group does not allow deletion.

HTTP Status Code: 400

### **InvalidDBSubnetGroupStateFault**

The DB subnet group cannot be deleted because it is in use.

HTTP Status Code: 400

### **InvalidSubnet**

The requested subnet is invalid, or multiple subnets were requested that are not all in a common VPC.

HTTP Status Code: 400

### **InvalidVPCNetworkStateFault**

DB subnet group does not cover all Availability Zones after it is created because users' change.

HTTP Status Code: 400

### **StorageQuotaExceeded**

Request would result in user exceeding the allowed amount of storage available across all DB instances.

HTTP Status Code: 400

### **StorageTypeNotSupported**

*StorageType* specified cannot be associated with the DB Instance.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ModifyDBClusterEndpoint

Modifies the properties of an endpoint in an Amazon Neptune DB cluster.

## Request Parameters

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

### DBClusterEndpointIdentifier

The identifier of the endpoint to modify. This parameter is stored as a lowercase string.

Type: String

Required: Yes

### EndpointType

The type of the endpoint. One of: READER, WRITER, ANY.

Type: String

Required: No

### ExcludedMembers.member.N

List of DB instance identifiers that aren't part of the custom endpoint group. All other eligible instances are reachable through the custom endpoint. Only relevant if the list of static members is empty.

Type: Array of strings

Required: No

### StaticMembers.member.N

List of DB instance identifiers that are part of the custom endpoint group.

Type: Array of strings

Required: No

## Response Elements

The following elements are returned by the service.

## **CustomEndpointType**

The type associated with a custom endpoint. One of: `READER`, `WRITER`, `ANY`.

Type: String

## **DBClusterEndpointArn**

The Amazon Resource Name (ARN) for the endpoint.

Type: String

## **DBClusterEndpointIdentifier**

The identifier associated with the endpoint. This parameter is stored as a lowercase string.

Type: String

## **DBClusterEndpointResourceIdentifier**

A unique system-generated identifier for an endpoint. It remains the same for the whole life of the endpoint.

Type: String

## **DBClusterIdentifier**

The DB cluster identifier of the DB cluster associated with the endpoint. This parameter is stored as a lowercase string.

Type: String

## **Endpoint**

The DNS address of the endpoint.

Type: String

## **EndpointType**

The type of the endpoint. One of: `READER`, `WRITER`, `CUSTOM`.

Type: String

## **ExcludedMembers.member.N**

List of DB instance identifiers that aren't part of the custom endpoint group. All other eligible instances are reachable through the custom endpoint. Only relevant if the list of static members is empty.

Type: Array of strings

### **StaticMembers.member.N**

List of DB instance identifiers that are part of the custom endpoint group.

Type: Array of strings

### **Status**

The current status of the endpoint. One of: `creating`, `available`, `deleting`, `inactive`, `modifying`. The `inactive` state applies to an endpoint that cannot be used for a certain kind of cluster, such as a `writer` endpoint for a read-only secondary cluster in a global database.

Type: String

## **Errors**

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

### **DBClusterEndpointNotFoundFault**

The specified custom endpoint doesn't exist.

HTTP Status Code: 400

### **DBInstanceNotFound**

*DBInstanceIdentifier* does not refer to an existing DB instance.

HTTP Status Code: 404

### **InvalidDBClusterEndpointStateFault**

The requested operation cannot be performed on the endpoint while the endpoint is in this state.

HTTP Status Code: 400

### **InvalidDBClusterStateFault**

The DB cluster is not in a valid state.

HTTP Status Code: 400

## InvalidDBInstanceState

The specified DB instance is not in the *available* state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ModifyDBClusterParameterGroup

Modifies the parameters of a DB cluster parameter group. To modify more than one parameter, submit a list of the following: `ParameterName`, `ParameterValue`, and `ApplyMethod`. A maximum of 20 parameters can be modified in a single request.

## Note

Changes to dynamic parameters are applied immediately. Changes to static parameters require a reboot without failover to the DB cluster associated with the parameter group before the change can take effect.

## Important

After you create a DB cluster parameter group, you should wait at least 5 minutes before creating your first DB cluster that uses that DB cluster parameter group as the default parameter group. This allows Amazon Neptune to fully complete the create action before the parameter group is used as the default for a new DB cluster. This is especially important for parameters that are critical when creating the default database for a DB cluster, such as the character set for the default database defined by the `character_set_database` parameter. You can use the *Parameter Groups* option of the Amazon Neptune console or the [DescribeDBClusterParameters](#) command to verify that your DB cluster parameter group has been created or modified.

## Request Parameters

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

### DBClusterParameterGroupName

The name of the DB cluster parameter group to modify.

Type: String

Required: Yes

## Parameters.Parameter.N

A list of parameters in the DB cluster parameter group to modify.

Type: Array of [Parameter](#) objects

Required: Yes

## Response Elements

The following element is returned by the service.

### DBClusterParameterGroupName

The name of the DB cluster parameter group.

Constraints:

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

#### Note

This value is stored as a lowercase string.

Type: String

## Errors

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

### DBParameterGroupNotFound

*DBParameterGroupName* does not refer to an existing DB parameter group.

HTTP Status Code: 404

## InvalidDBParameterGroupState

The DB parameter group is in use or is in an invalid state. If you are attempting to delete the parameter group, you cannot delete it when the parameter group is in this state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# ModifyDBClusterSnapshotAttribute

Adds an attribute and values to, or removes an attribute and values from, a manual DB cluster snapshot.

To share a manual DB cluster snapshot with other Amazon accounts, specify `restore` as the `AttributeName` and use the `ValuesToAdd` parameter to add a list of IDs of the Amazon accounts that are authorized to restore the manual DB cluster snapshot. Use the value `all` to make the manual DB cluster snapshot public, which means that it can be copied or restored by all Amazon accounts. Do not add the `all` value for any manual DB cluster snapshots that contain private information that you don't want available to all Amazon accounts. If a manual DB cluster snapshot is encrypted, it can be shared, but only by specifying a list of authorized Amazon account IDs for the `ValuesToAdd` parameter. You can't use `all` as a value for that parameter in this case.

To view which Amazon accounts have access to copy or restore a manual DB cluster snapshot, or whether a manual DB cluster snapshot public or private, use the [DescribeDBClusterSnapshotAttributes](#) API action.

## Request Parameters

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

### AttributeName

The name of the DB cluster snapshot attribute to modify.

To manage authorization for other Amazon accounts to copy or restore a manual DB cluster snapshot, set this value to `restore`.

Type: String

Required: Yes

### DBClusterSnapshotIdentifier

The identifier for the DB cluster snapshot to modify the attributes for.

Type: String

Required: Yes

## ValuesToAdd.AttributeValue.N

A list of DB cluster snapshot attributes to add to the attribute specified by `AttributeName`.

To authorize other Amazon accounts to copy or restore a manual DB cluster snapshot, set this list to include one or more Amazon account IDs, or `all` to make the manual DB cluster snapshot restorable by any Amazon account. Do not add the `all` value for any manual DB cluster snapshots that contain private information that you don't want available to all Amazon accounts.

Type: Array of strings

Required: No

## ValuesToRemove.AttributeValue.N

A list of DB cluster snapshot attributes to remove from the attribute specified by `AttributeName`.

To remove authorization for other Amazon accounts to copy or restore a manual DB cluster snapshot, set this list to include one or more Amazon account identifiers, or `all` to remove authorization for any Amazon account to copy or restore the DB cluster snapshot. If you specify `all`, an Amazon account whose account ID is explicitly added to the `restore` attribute can still copy or restore a manual DB cluster snapshot.

Type: Array of strings

Required: No

## Response Elements

The following element is returned by the service.

### DBClusterSnapshotAttributesResult

Contains the results of a successful call to the [DescribeDBClusterSnapshotAttributes](#) API action.

Manual DB cluster snapshot attributes are used to authorize other Amazon accounts to copy or restore a manual DB cluster snapshot. For more information, see the [ModifyDBClusterSnapshotAttribute](#) API action.

Type: [DBClusterSnapshotAttributesResult](#) object

## Errors

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

### **DBClusterSnapshotNotFoundFault**

*DBClusterSnapshotIdentifier* does not refer to an existing DB cluster snapshot.

HTTP Status Code: 404

### **InvalidDBClusterSnapshotStateFault**

The supplied value is not a valid DB cluster snapshot state.

HTTP Status Code: 400

### **SharedSnapshotQuotaExceeded**

You have exceeded the maximum number of accounts that you can share a manual DB snapshot with.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ModifyDBInstance

Modifies settings for a DB instance. You can change one or more database configuration parameters by specifying these parameters and the new values in the request. To learn what modifications you can make to your DB instance, call [DescribeValidDBInstanceModifications](#) before you call [ModifyDBInstance](#).

## Request Parameters

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

### AllocatedStorage

Not supported by Neptune.

Type: Integer

Required: No

### AllowMajorVersionUpgrade

Indicates that major version upgrades are allowed. Changing this parameter doesn't result in an outage and the change is asynchronously applied as soon as possible.

Type: Boolean

Required: No

### ApplyImmediately

Specifies whether the modifications in this request and any pending modifications are asynchronously applied as soon as possible, regardless of the PreferredMaintenanceWindow setting for the DB instance.

If this parameter is set to `false`, changes to the DB instance are applied during the next maintenance window. Some parameter changes can cause an outage and are applied on the next call to [RebootDBInstance](#), or the next failure reboot.

Default: `false`

Type: Boolean

Required: No

## AutoMinorVersionUpgrade

Indicates that minor version upgrades are applied automatically to the DB instance during the maintenance window. Changing this parameter doesn't result in an outage except in the following case and the change is asynchronously applied as soon as possible. An outage will result if this parameter is set to `true` during the maintenance window, and a newer minor version is available, and Neptune has enabled auto patching for that engine version.

Type: Boolean

Required: No

## BackupRetentionPeriod

Not applicable. The retention period for automated backups is managed by the DB cluster. For more information, see [ModifyDBCluster](#).

Default: Uses existing setting

Type: Integer

Required: No

## CACertificateIdentifier

Indicates the certificate that needs to be associated with the instance.

Type: String

Required: No

## CloudwatchLogsExportConfiguration

The configuration setting for the log types to be enabled for export to CloudWatch Logs for a specific DB instance or DB cluster.

Type: [CloudwatchLogsExportConfiguration](#) object

Required: No

## CopyTagsToSnapshot

True to copy all tags from the DB instance to snapshots of the DB instance, and otherwise false. The default is false.

Type: Boolean

Required: No

### **DBInstanceClass**

The new compute and memory capacity of the DB instance, for example, `db.m4.large`. Not all DB instance classes are available in all Amazon Regions.

If you modify the DB instance class, an outage occurs during the change. The change is applied during the next maintenance window, unless `ApplyImmediately` is specified as `true` for this request.

Default: Uses existing setting

Type: String

Required: No

### **DBInstanceIdentifier**

The DB instance identifier. This value is stored as a lowercase string.

Constraints:

- Must match the identifier of an existing DBInstance.

Type: String

Required: Yes

### **DBParameterGroupName**

The name of the DB parameter group to apply to the DB instance. Changing this setting doesn't result in an outage. The parameter group name itself is changed immediately, but the actual parameter changes are not applied until you reboot the instance without failover. The db instance will NOT be rebooted automatically and the parameter changes will NOT be applied during the next maintenance window.

Default: Uses existing setting

Constraints: The DB parameter group must be in the same DB parameter group family as this DB instance.

Type: String

Required: No

### **DBPortNumber**

The port number on which the database accepts connections.

The value of the `DBPortNumber` parameter must not match any of the port values specified for options in the option group for the DB instance.

Your database will restart when you change the `DBPortNumber` value regardless of the value of the `ApplyImmediately` parameter.

Default: 8182

Type: Integer

Required: No

### **DBSecurityGroups.DBSecurityGroupName.N**

A list of DB security groups to authorize on this DB instance. Changing this setting doesn't result in an outage and the change is asynchronously applied as soon as possible.

Constraints:

- If supplied, must match existing `DBSecurityGroups`.

Type: Array of strings

Required: No

### **DBSubnetGroupName**

The new DB subnet group for the DB instance. You can use this parameter to move your DB instance to a different VPC.

Changing the subnet group causes an outage during the change. The change is applied during the next maintenance window, unless you specify `true` for the `ApplyImmediately` parameter.

Constraints: If supplied, must match the name of an existing `DBSubnetGroup`.

Example: `mySubnetGroup`

Type: String

Required: No

## DeletionProtection

A value that indicates whether the DB instance has deletion protection enabled. The database can't be deleted when deletion protection is enabled. By default, deletion protection is disabled. See [Deleting a DB Instance](#).

Type: Boolean

Required: No

## Domain

Not supported.

Type: String

Required: No

## DomainIAMRoleName

Not supported

Type: String

Required: No

## EnableIAMDatabaseAuthentication

True to enable mapping of Amazon Identity and Access Management (IAM) accounts to database accounts, and otherwise false.

You can enable IAM database authentication for the following database engines

Not applicable. Mapping Amazon IAM accounts to database accounts is managed by the DB cluster. For more information, see [ModifyDBCluster](#).

Default: false

Type: Boolean

Required: No

## EnablePerformanceInsights

*(Not supported by Neptune)*



Type: Boolean

Required: No

### **EngineVersion**

The version number of the database engine to upgrade to. Currently, setting this parameter has no effect. To upgrade your database engine to the most recent release, use the [ApplyPendingMaintenanceAction](#) API.

Type: String

Required: No

### **Iops**

The new Provisioned IOPS (I/O operations per second) value for the instance.

Changing this setting doesn't result in an outage and the change is applied during the next maintenance window unless the `ApplyImmediately` parameter is set to `true` for this request.

Default: Uses existing setting

Type: Integer

Required: No

### **LicenseModel**

Not supported by Neptune.

Type: String

Required: No

### **MasterUserPassword**

Not supported by Neptune.

Type: String

Required: No

### **MonitoringInterval**

The interval, in seconds, between points when Enhanced Monitoring metrics are collected for the DB instance. To disable collecting Enhanced Monitoring metrics, specify 0. The default is 0.

If `MonitoringRoleArn` is specified, then you must also set `MonitoringInterval` to a value other than 0.

Valid Values: 0, 1, 5, 10, 15, 30, 60

Type: Integer

Required: No

### **MonitoringRoleArn**

The ARN for the IAM role that permits Neptune to send enhanced monitoring metrics to Amazon CloudWatch Logs. For example, `arn:aws:iam:123456789012:role/emaccess`.

If `MonitoringInterval` is set to a value other than 0, then you must supply a `MonitoringRoleArn` value.

Type: String

Required: No

### **MultiAZ**

Specifies if the DB instance is a Multi-AZ deployment. Changing this parameter doesn't result in an outage and the change is applied during the next maintenance window unless the `ApplyImmediately` parameter is set to `true` for this request.

Type: Boolean

Required: No

### **NewDBInstanceIdentifier**

The new DB instance identifier for the DB instance when renaming a DB instance. When you change the DB instance identifier, an instance reboot will occur immediately if you set `Apply Immediately` to `true`, or will occur during the next maintenance window if `Apply Immediately` to `false`. This value is stored as a lowercase string.

Constraints:

- Must contain from 1 to 63 letters, numbers, or hyphens.
- The first character must be a letter.
- Cannot end with a hyphen or contain two consecutive hyphens.

Example: mydbinstance

Type: String

Required: No

### **OptionGroupName**

*(Not supported by Neptune)*

Type: String

Required: No

### **PerformanceInsightsKMSKeyId**

*(Not supported by Neptune)*

Type: String

Required: No

### **PreferredBackupWindow**

The daily time range during which automated backups are created if automated backups are enabled.

Not applicable. The daily time range for creating automated backups is managed by the DB cluster. For more information, see [ModifyDBCluster](#).

Constraints:

- Must be in the format hh24:mi-hh24:mi
- Must be in Universal Time Coordinated (UTC)
- Must not conflict with the preferred maintenance window
- Must be at least 30 minutes

Type: String

Required: No

### **PreferredMaintenanceWindow**

The weekly time range (in UTC) during which system maintenance can occur, which might result in an outage. Changing this parameter doesn't result in an outage, except in the following

situation, and the change is asynchronously applied as soon as possible. If there are pending actions that cause a reboot, and the maintenance window is changed to include the current time, then changing this parameter will cause a reboot of the DB instance. If moving this window to the current time, there must be at least 30 minutes between the current time and end of the window to ensure pending changes are applied.

Default: Uses existing setting

Format: ddd:hh24:mi-ddd:hh24:mi

Valid Days: Mon | Tue | Wed | Thu | Fri | Sat | Sun

Constraints: Must be at least 30 minutes

Type: String

Required: No

### **PromotionTier**

A value that specifies the order in which a Read Replica is promoted to the primary instance after a failure of the existing primary instance.

Default: 1

Valid Values: 0 - 15

Type: Integer

Required: No

### **PubliclyAccessible**

*This parameter has been deprecated.*

This flag should no longer be used.

Type: Boolean

Required: No

### **StorageType**

Not applicable. In Neptune the storage type is managed at the DB Cluster level.

Type: String

Required: No

### **TdeCredentialArn**

The ARN from the key store with which to associate the instance for TDE encryption.

Type: String

Required: No

### **TdeCredentialPassword**

The password for the given ARN from the key store in order to access the device.

Type: String

Required: No

### **VpcSecurityGroupIds.VpcSecurityGroupId.N**

A list of EC2 VPC security groups to authorize on this DB instance. This change is asynchronously applied as soon as possible.

Not applicable. The associated list of EC2 VPC security groups is managed by the DB cluster. For more information, see [ModifyDBCluster](#).

Constraints:

- If supplied, must match existing VpcSecurityGroupIds.

Type: Array of strings

Required: No

## **Response Elements**

The following element is returned by the service.

### **DBInstance**

Contains the details of an Amazon Neptune DB instance.

This data type is used as a response element in the [DescribeDBInstances](#) action.

Type: [DBInstance](#) object

## Errors

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

### AuthorizationNotFound

Specified CIDRIP or EC2 security group is not authorized for the specified DB security group.

Neptune may not also be authorized via IAM to perform necessary actions on your behalf.

HTTP Status Code: 404

### CertificateNotFound

*CertificateIdentifier* does not refer to an existing certificate.

HTTP Status Code: 404

### DBInstanceAlreadyExists

User already has a DB instance with the given identifier.

HTTP Status Code: 400

### DBInstanceNotFound

*DBInstanceIdentifier* does not refer to an existing DB instance.

HTTP Status Code: 404

### DBParameterGroupNotFound

*DBParameterGroupName* does not refer to an existing DB parameter group.

HTTP Status Code: 404

### DBSecurityGroupNotFound

*DBSecurityGroupName* does not refer to an existing DB security group.

HTTP Status Code: 404

### DBUpgradeDependencyFailure

The DB upgrade failed because a resource the DB depends on could not be modified.

HTTP Status Code: 400

### **DomainNotFoundFault**

*Domain* does not refer to an existing Active Directory Domain.

HTTP Status Code: 404

### **InsufficientDBInstanceCapacity**

Specified DB instance class is not available in the specified Availability Zone.

HTTP Status Code: 400

### **InvalidDBInstanceState**

The specified DB instance is not in the *available* state.

HTTP Status Code: 400

### **InvalidDBSecurityGroupState**

The state of the DB security group does not allow deletion.

HTTP Status Code: 400

### **InvalidVPCNetworkStateFault**

DB subnet group does not cover all Availability Zones after it is created because users' change.

HTTP Status Code: 400

### **OptionGroupNotFoundFault**

The designated option group could not be found.

HTTP Status Code: 404

### **ProvisionedIopsNotAvailableInAZFault**

Provisioned IOPS not available in the specified Availability Zone.

HTTP Status Code: 400

### **StorageQuotaExceeded**

Request would result in user exceeding the allowed amount of storage available across all DB instances.

HTTP Status Code: 400

### **StorageTypeNotSupported**

*StorageType* specified cannot be associated with the DB Instance.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# ModifyDBParameterGroup

Modifies the parameters of a DB parameter group. To modify more than one parameter, submit a list of the following: `ParameterName`, `ParameterValue`, and `ApplyMethod`. A maximum of 20 parameters can be modified in a single request.

## Note

Changes to dynamic parameters are applied immediately. Changes to static parameters require a reboot without failover to the DB instance associated with the parameter group before the change can take effect.

## Important

After you modify a DB parameter group, you should wait at least 5 minutes before creating your first DB instance that uses that DB parameter group as the default parameter group. This allows Amazon Neptune to fully complete the modify action before the parameter group is used as the default for a new DB instance. This is especially important for parameters that are critical when creating the default database for a DB instance, such as the character set for the default database defined by the `character_set_database` parameter. You can use the *Parameter Groups* option of the Amazon Neptune console or the *DescribeDBParameters* command to verify that your DB parameter group has been created or modified.

## Request Parameters

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

### DBParameterGroupName

The name of the DB parameter group.

Constraints:

- If supplied, must match the name of an existing `DBParameterGroup`.

Type: String

Required: Yes

### Parameters.Parameter.N

An array of parameter names, values, and the apply method for the parameter update. At least one parameter name, value, and apply method must be supplied; subsequent arguments are optional. A maximum of 20 parameters can be modified in a single request.

Valid Values (for the application method): `immediate` | `pending-reboot`

#### Note

You can use the `immediate` value with dynamic parameters only. You can use the `pending-reboot` value for both dynamic and static parameters, and changes are applied when you reboot the DB instance without failover.

Type: Array of [Parameter](#) objects

Required: Yes

## Response Elements

The following element is returned by the service.

### DBParameterGroupName

Provides the name of the DB parameter group.

Type: String

## Errors

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

### DBParameterGroupNotFound

*DBParameterGroupName* does not refer to an existing DB parameter group.

HTTP Status Code: 404

## InvalidDBParameterGroupState

The DB parameter group is in use or is in an invalid state. If you are attempting to delete the parameter group, you cannot delete it when the parameter group is in this state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ModifyDBSubnetGroup

Modifies an existing DB subnet group. DB subnet groups must contain at least one subnet in at least two AZs in the Amazon Region.

## Request Parameters

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

### DBSubnetGroupDescription

The description for the DB subnet group.

Type: String

Required: No

### DBSubnetGroupName

The name for the DB subnet group. This value is stored as a lowercase string. You can't modify the default subnet group.

Constraints: Must match the name of an existing DBSubnetGroup. Must not be default.

Example: mySubnetgroup

Type: String

Required: Yes

### SubnetIds.SubnetIdentifier.N

The EC2 subnet IDs for the DB subnet group.

Type: Array of strings

Required: Yes

## Response Elements

The following element is returned by the service.

## DBSubnetGroup

Contains the details of an Amazon Neptune DB subnet group.

This data type is used as a response element in the [DescribeDBSubnetGroups](#) action.

Type: [DBSubnetGroup](#) object

## Errors

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

### DBSubnetGroupDoesNotCoverEnoughAZs

Subnets in the DB subnet group should cover at least two Availability Zones unless there is only one Availability Zone.

HTTP Status Code: 400

### DBSubnetGroupNotFoundFault

*DBSubnetGroupName* does not refer to an existing DB subnet group.

HTTP Status Code: 404

### DBSubnetQuotaExceededFault

Request would result in user exceeding the allowed number of subnets in a DB subnet groups.

HTTP Status Code: 400

### InvalidSubnet

The requested subnet is invalid, or multiple subnets were requested that are not all in a common VPC.

HTTP Status Code: 400

### SubnetAlreadyInUse

The DB subnet is already in use in the Availability Zone.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ModifyEventSubscription

Modifies an existing event notification subscription. Note that you can't modify the source identifiers using this call; to change source identifiers for a subscription, use the [AddSourceIdentifierToSubscription](#) and [RemoveSourceIdentifierFromSubscription](#) calls.

You can see a list of the event categories for a given `SourceType` by using the **DescribeEventCategories** action.

## Request Parameters

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

### Enabled

A Boolean value; set to **true** to activate the subscription.

Type: Boolean

Required: No

### EventCategories.EventCategory.N

A list of event categories for a `SourceType` that you want to subscribe to. You can see a list of the categories for a given `SourceType` by using the **DescribeEventCategories** action.

Type: Array of strings

Required: No

### SnsTopicArn

The Amazon Resource Name (ARN) of the SNS topic created for event notification. The ARN is created by Amazon SNS when you create a topic and subscribe to it.

Type: String

Required: No

### SourceType

The type of source that is generating the events. For example, if you want to be notified of events generated by a DB instance, you would set this parameter to `db-instance`. If this value is not specified, all events are returned.

Valid values: db-instance | db-parameter-group | db-security-group | db-snapshot

Type: String

Required: No

### **SubscriptionName**

The name of the event notification subscription.

Type: String

Required: Yes

## **Response Elements**

The following element is returned by the service.

### **EventSubscription**

Contains the results of a successful invocation of the [DescribeEventSubscriptions](#) action.

Type: [EventSubscription](#) object

## **Errors**

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

### **EventSubscriptionQuotaExceeded**

You have exceeded the number of events you can subscribe to.

HTTP Status Code: 400

### **SNSInvalidTopic**

The SNS topic is invalid.

HTTP Status Code: 400

### **SNSNoAuthorization**

There is no SNS authorization.



HTTP Status Code: 400

### **SNSTopicArnNotFound**

The ARN of the SNS topic could not be found.

HTTP Status Code: 404

### **SubscriptionCategoryNotFound**

The designated subscription category could not be found.

HTTP Status Code: 404

### **SubscriptionNotFound**

The designated subscription 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ModifyGlobalCluster

Modify a setting for an Amazon Neptune global cluster. You can change one or more database configuration parameters by specifying these parameters and their new values in the request.

## Request Parameters

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

### AllowMajorVersionUpgrade

A value that indicates whether major version upgrades are allowed.

**Constraints:** You must allow major version upgrades if you specify a value for the `EngineVersion` parameter that is a different major version than the DB cluster's current version.

If you upgrade the major version of a global database, the cluster and DB instance parameter groups are set to the default parameter groups for the new version, so you will need to apply any custom parameter groups after completing the upgrade.

Type: Boolean

Required: No

### DeletionProtection

Indicates whether the global database has deletion protection enabled. The global database cannot be deleted when deletion protection is enabled.

Type: Boolean

Required: No

### EngineVersion

The version number of the database engine to which you want to upgrade. Changing this parameter will result in an outage. The change is applied during the next maintenance window unless `ApplyImmediately` is enabled.

To list all of the available Neptune engine versions, use the following command:

Type: String

Required: No

### **GlobalClusterIdentifier**

The DB cluster identifier for the global cluster being modified. This parameter is not case-sensitive.

Constraints: Must match the identifier of an existing global database cluster.

Type: String

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

Pattern: `[A-Za-z][0-9A-Za-z-:._]*`

Required: Yes

### **NewGlobalClusterIdentifier**

A new cluster identifier to assign to the global database. This value is stored as a lowercase string.

Constraints:

- Must contain from 1 to 63 letters, numbers, or hyphens.
- The first character must be a letter.
- Can't end with a hyphen or contain two consecutive hyphens

Example: `my-cluster2`

Type: String

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

Pattern: `[A-Za-z][0-9A-Za-z-:._]*`

Required: No

## **Response Elements**

The following element is returned by the service.

## GlobalCluster

Contains the details of an Amazon Neptune global database.

This data type is used as a response element for the [CreateGlobalCluster](#), [DescribeGlobalClusters](#), [ModifyGlobalCluster](#), [DeleteGlobalCluster](#), [FailoverGlobalCluster](#), and [RemoveFromGlobalCluster](#) actions.

Type: [GlobalCluster](#) object

## Errors

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

### GlobalClusterNotFoundFault

The `GlobalClusterIdentifier` doesn't refer to an existing global database cluster.

HTTP Status Code: 404

### InvalidGlobalClusterStateFault

The global cluster is in an invalid state and can't perform the requested operation.

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 PHP V3](#)
- [AWS SDK for Python](#)

- [AWS SDK for Ruby V3](#)

# PromoteReadReplicaDBCluster

Not supported.

## Request Parameters

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

### DBClusterIdentifier

Not supported.

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### DBCluster

Contains the details of an Amazon Neptune DB cluster.

This data type is used as a response element in the [DescribeDBClusters](#).

Type: [DBCluster](#) object

## Errors

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

### DBClusterNotFoundFault

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### InvalidDBClusterStateFault

The DB cluster is not in a valid state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# RebootDBInstance

You might need to reboot your DB instance, usually for maintenance reasons. For example, if you make certain modifications, or if you change the DB parameter group associated with the DB instance, you must reboot the instance for the changes to take effect.

Rebooting a DB instance restarts the database engine service. Rebooting a DB instance results in a momentary outage, during which the DB instance status is set to rebooting.

## Request Parameters

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

### DBInstanceIdentifier

The DB instance identifier. This parameter is stored as a lowercase string.

Constraints:

- Must match the identifier of an existing DBInstance.

Type: String

Required: Yes

### ForceFailover

When `true`, the reboot is conducted through a MultiAZ failover.

Constraint: You can't specify `true` if the instance is not configured for MultiAZ.

Type: Boolean

Required: No

## Response Elements

The following element is returned by the service.

### DBInstance

Contains the details of an Amazon Neptune DB instance.



This data type is used as a response element in the [DescribeDBInstances](#) action.

Type: [DBInstance](#) object

## Errors

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

### DBInstanceNotFound

*DBInstanceIdentifier* does not refer to an existing DB instance.

HTTP Status Code: 404

### InvalidDBInstanceState

The specified DB instance is not in the *available* state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# RemoveFromGlobalCluster

Detaches a Neptune DB cluster from a Neptune global database. A secondary cluster becomes a normal standalone cluster with read-write capability instead of being read-only, and no longer receives data from a the primary cluster.

## Request Parameters

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

### DbClusterIdentifier

The Amazon Resource Name (ARN) identifying the cluster to be detached from the Neptune global database cluster.

Type: String

Required: Yes

### GlobalClusterIdentifier

The identifier of the Neptune global database from which to detach the specified Neptune DB cluster.

Type: String

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

Pattern: `[A-Za-z][0-9A-Za-z-:._]*`

Required: Yes

## Response Elements

The following element is returned by the service.

### GlobalCluster

Contains the details of an Amazon Neptune global database.

This data type is used as a response element for the [CreateGlobalCluster](#), [DescribeGlobalClusters](#), [ModifyGlobalCluster](#), [DeleteGlobalCluster](#), [FailoverGlobalCluster](#), and [RemoveFromGlobalCluster](#) actions.

Type: [GlobalCluster](#) object

## Errors

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

### **DBClusterNotFoundFault**

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### **GlobalClusterNotFoundFault**

The *GlobalClusterIdentifier* doesn't refer to an existing global database cluster.

HTTP Status Code: 404

### **InvalidGlobalClusterStateFault**

The global cluster is in an invalid state and can't perform the requested operation.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# RemoveRoleFromDBCluster

Disassociates an Identity and Access Management (IAM) role from a DB cluster.

## Request Parameters

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

### DBClusterIdentifier

The name of the DB cluster to disassociate the IAM role from.

Type: String

Required: Yes

### FeatureName

The name of the feature for the DB cluster that the IAM role is to be disassociated from. For the list of supported feature names, see [DescribeDBEngineVersions](#).

Type: String

Required: No

### RoleArn

The Amazon Resource Name (ARN) of the IAM role to disassociate from the DB cluster, for example `arn:aws:iam::123456789012:role/NeptuneAccessRole`.

Type: String

Required: Yes

## Errors

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

### DBClusterNotFoundFault

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

## **DBClusterRoleNotFound**

The specified IAM role Amazon Resource Name (ARN) is not associated with the specified DB cluster.

HTTP Status Code: 404

## **InvalidDBClusterStateFault**

The DB cluster is not in a valid state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# RemoveSourceIdentifierFromSubscription

Removes a source identifier from an existing event notification subscription.

## Request Parameters

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

### SourceIdentifier

The source identifier to be removed from the subscription, such as the **DB instance identifier** for a DB instance or the name of a security group.

Type: String

Required: Yes

### SubscriptionName

The name of the event notification subscription you want to remove a source identifier from.

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### EventSubscription

Contains the results of a successful invocation of the [DescribeEventSubscriptions](#) action.

Type: [EventSubscription](#) object

## Errors

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

### SourceNotFound

The source could not be found.

HTTP Status Code: 404

## SubscriptionNotFound

The designated subscription 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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# RemoveTagsFromResource

Removes metadata tags from an Amazon Neptune resource.

## Request Parameters

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

### ResourceName

The Amazon Neptune resource that the tags are removed from. This value is an Amazon Resource Name (ARN). For information about creating an ARN, see [Constructing an Amazon Resource Name \(ARN\)](#).

Type: String

Required: Yes

### TagKeys.member.N

The tag key (name) of the tag to be removed.

Type: Array of strings

Required: Yes

## Errors

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

### DBClusterNotFoundFault

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### DBInstanceNotFound

*DBInstanceIdentifier* does not refer to an existing DB instance.

HTTP Status Code: 404



## DBSnapshotNotFound

*DBSnapshotIdentifier* does not refer to an existing DB snapshot.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ResetDBClusterParameterGroup

Modifies the parameters of a DB cluster parameter group to the default value. To reset specific parameters submit a list of the following: `ParameterName` and `ApplyMethod`. To reset the entire DB cluster parameter group, specify the `DBClusterParameterGroupName` and `ResetAllParameters` parameters.

When resetting the entire group, dynamic parameters are updated immediately and static parameters are set to `pending-reboot` to take effect on the next DB instance restart or [RebootDBInstance](#) request. You must call [RebootDBInstance](#) for every DB instance in your DB cluster that you want the updated static parameter to apply to.

## Request Parameters

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

### **DBClusterParameterGroupName**

The name of the DB cluster parameter group to reset.

Type: String

Required: Yes

### **Parameters.Parameter.N**

A list of parameter names in the DB cluster parameter group to reset to the default values. You can't use this parameter if the `ResetAllParameters` parameter is set to `true`.

Type: Array of [Parameter](#) objects

Required: No

### **ResetAllParameters**

A value that is set to `true` to reset all parameters in the DB cluster parameter group to their default values, and `false` otherwise. You can't use this parameter if there is a list of parameter names specified for the `Parameters` parameter.

Type: Boolean

Required: No

## Response Elements

The following element is returned by the service.

### **DBClusterParameterGroupName**

The name of the DB cluster parameter group.

Constraints:

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

 **Note**

This value is stored as a lowercase string.

Type: String

## Errors

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

### **DBParameterGroupNotFound**

*DBParameterGroupName* does not refer to an existing DB parameter group.

HTTP Status Code: 404

### **InvalidDBParameterGroupState**

The DB parameter group is in use or is in an invalid state. If you are attempting to delete the parameter group, you cannot delete it when the parameter group is in this state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ResetDBParameterGroup

Modifies the parameters of a DB parameter group to the engine/system default value. To reset specific parameters, provide a list of the following: `ParameterName` and `ApplyMethod`. To reset the entire DB parameter group, specify the `DBParameterGroup` name and `ResetAllParameters` parameters. When resetting the entire group, dynamic parameters are updated immediately and static parameters are set to `pending-reboot` to take effect on the next DB instance restart or `RebootDBInstance` request.

## Request Parameters

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

### DBParameterGroupName

The name of the DB parameter group.

Constraints:

- Must match the name of an existing `DBParameterGroup`.

Type: String

Required: Yes

### Parameters.Parameter.N

To reset the entire DB parameter group, specify the `DBParameterGroup` name and `ResetAllParameters` parameters. To reset specific parameters, provide a list of the following: `ParameterName` and `ApplyMethod`. A maximum of 20 parameters can be modified in a single request.

Valid Values (for Apply method): `pending-reboot`

Type: Array of [Parameter](#) objects

Required: No

### ResetAllParameters

Specifies whether (`true`) or not (`false`) to reset all parameters in the DB parameter group to default values.

Default: true

Type: Boolean

Required: No

## Response Elements

The following element is returned by the service.

### **DBParameterGroupName**

Provides the name of the DB parameter group.

Type: String

## Errors

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

### **DBParameterGroupNotFound**

*DBParameterGroupName* does not refer to an existing DB parameter group.

HTTP Status Code: 404

### **InvalidDBParameterGroupState**

The DB parameter group is in use or is in an invalid state. If you are attempting to delete the parameter group, you cannot delete it when the parameter group is in this state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# RestoreDBClusterFromSnapshot

Creates a new DB cluster from a DB snapshot or DB cluster snapshot.

If a DB snapshot is specified, the target DB cluster is created from the source DB snapshot with a default configuration and default security group.

If a DB cluster snapshot is specified, the target DB cluster is created from the source DB cluster restore point with the same configuration as the original source DB cluster, except that the new DB cluster is created with the default security group.

## Request Parameters

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

### AvailabilityZones.AvailabilityZone.N

Provides the list of EC2 Availability Zones that instances in the restored DB cluster can be created in.

Type: Array of strings

Required: No

### CopyTagsToSnapshot

*If set to `true`, tags are copied to any snapshot of the restored DB cluster that is created.*

Type: Boolean

Required: No

### DatabaseName

Not supported.

Type: String

Required: No

### DBClusterIdentifier

The name of the DB cluster to create from the DB snapshot or DB cluster snapshot. This parameter isn't case-sensitive.



**Constraints:**

- Must contain from 1 to 63 letters, numbers, or hyphens
- First character must be a letter
- Cannot end with a hyphen or contain two consecutive hyphens

Example: my-snapshot-id

Type: String

Required: Yes

**DBClusterParameterGroupName**

The name of the DB cluster parameter group to associate with the new DB cluster.

**Constraints:**

- If supplied, must match the name of an existing DBClusterParameterGroup.

Type: String

Required: No

**DBSubnetGroupName**

The name of the DB subnet group to use for the new DB cluster.

Constraints: If supplied, must match the name of an existing DBSubnetGroup.

Example: mySubnetgroup

Type: String

Required: No

**DeletionProtection**

A value that indicates whether the DB cluster has deletion protection enabled. The database can't be deleted when deletion protection is enabled. By default, deletion protection is disabled.

Type: Boolean

Required: No

## **EnableCloudwatchLogsExports.member.N**

The list of logs that the restored DB cluster is to export to Amazon CloudWatch Logs.

Type: Array of strings

Required: No

## **EnableIAMDatabaseAuthentication**

True to enable mapping of Amazon Identity and Access Management (IAM) accounts to database accounts, and otherwise false.

Default: false

Type: Boolean

Required: No

## **Engine**

The database engine to use for the new DB cluster.

Default: The same as source

Constraint: Must be compatible with the engine of the source

Type: String

Required: Yes

## **EngineVersion**

The version of the database engine to use for the new DB cluster.

Type: String

Required: No

## **KmsKeyId**

The Amazon KMS key identifier to use when restoring an encrypted DB cluster from a DB snapshot or DB cluster snapshot.

The KMS key identifier is the Amazon Resource Name (ARN) for the KMS encryption key. If you are restoring a DB cluster with the same Amazon account that owns the KMS encryption key

used to encrypt the new DB cluster, then you can use the KMS key alias instead of the ARN for the KMS encryption key.

If you do not specify a value for the `KmsKeyId` parameter, then the following will occur:

- If the DB snapshot or DB cluster snapshot in `SnapshotIdentifier` is encrypted, then the restored DB cluster is encrypted using the KMS key that was used to encrypt the DB snapshot or DB cluster snapshot.
- If the DB snapshot or DB cluster snapshot in `SnapshotIdentifier` is not encrypted, then the restored DB cluster is not encrypted.

Type: String

Required: No

### **OptionGroupName**

*(Not supported by Neptune)*

Type: String

Required: No

### **Port**

The port number on which the new DB cluster accepts connections.

Constraints: Value must be 1150-65535

Default: The same port as the original DB cluster.

Type: Integer

Required: No

### **ServerlessV2ScalingConfiguration**

Contains the scaling configuration of a Neptune Serverless DB cluster.

For more information, see [Using Amazon Neptune Serverless](#) in the *Amazon Neptune User Guide*.

Type: [ServerlessV2ScalingConfiguration](#) object

Required: No

## SnapshotIdentifier

The identifier for the DB snapshot or DB cluster snapshot to restore from.

You can use either the name or the Amazon Resource Name (ARN) to specify a DB cluster snapshot. However, you can use only the ARN to specify a DB snapshot.

Constraints:

- Must match the identifier of an existing Snapshot.

Type: String

Required: Yes

## StorageType

Specifies the storage type to be associated with the DB cluster.

Valid values: `standard`, `iopt1`

Default: `standard`

Type: String

Required: No

## Tags.Tag.N

The tags to be assigned to the restored DB cluster.

Type: Array of [Tag](#) objects

Required: No

## VpcSecurityGroupIds.VpcSecurityGroupId.N

A list of VPC security groups that the new DB cluster will belong to.

Type: Array of strings

Required: No

## Response Elements

The following element is returned by the service.

## DBCluster

Contains the details of an Amazon Neptune DB cluster.

This data type is used as a response element in the [DescribeDBClusters](#).

Type: [DBCluster](#) object

## Errors

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

### DBClusterAlreadyExistsFault

User already has a DB cluster with the given identifier.

HTTP Status Code: 400

### DBClusterParameterGroupNotFound

*DBClusterParameterGroupName* does not refer to an existing DB Cluster parameter group.

HTTP Status Code: 404

### DBClusterQuotaExceededFault

User attempted to create a new DB cluster and the user has already reached the maximum allowed DB cluster quota.

HTTP Status Code: 403

### DBClusterSnapshotNotFoundFault

*DBClusterSnapshotIdentifier* does not refer to an existing DB cluster snapshot.

HTTP Status Code: 404

### DBSnapshotNotFound

*DBSnapshotIdentifier* does not refer to an existing DB snapshot.

HTTP Status Code: 404

### DBSubnetGroupNotFoundFault

*DBSubnetGroupName* does not refer to an existing DB subnet group.

HTTP Status Code: 404

### **DBSubnetGroupNotFoundFault**

*DBSubnetGroupName* does not refer to an existing DB subnet group.

HTTP Status Code: 404

### **InsufficientDBClusterCapacityFault**

The DB cluster does not have enough capacity for the current operation.

HTTP Status Code: 403

### **InsufficientStorageClusterCapacity**

There is insufficient storage available for the current action. You may be able to resolve this error by updating your subnet group to use different Availability Zones that have more storage available.

HTTP Status Code: 400

### **InvalidDBClusterSnapshotStateFault**

The supplied value is not a valid DB cluster snapshot state.

HTTP Status Code: 400

### **InvalidDBSnapshotState**

The state of the DB snapshot does not allow deletion.

HTTP Status Code: 400

### **InvalidRestoreFault**

Cannot restore from vpc backup to non-vpc DB instance.

HTTP Status Code: 400

### **InvalidSubnet**

The requested subnet is invalid, or multiple subnets were requested that are not all in a common VPC.

HTTP Status Code: 400

## InvalidVPCNetworkStateFault

DB subnet group does not cover all Availability Zones after it is created because users' change.

HTTP Status Code: 400

## KMSKeyNotAccessibleFault

Error accessing KMS key.

HTTP Status Code: 400

## OptionGroupNotFoundFault

The designated option group could not be found.

HTTP Status Code: 404

## StorageQuotaExceeded

Request would result in user exceeding the allowed amount of storage available across all DB instances.

HTTP Status Code: 400

## StorageQuotaExceeded

Request would result in user exceeding the allowed amount of storage available across all DB instances.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# RestoreDBClusterToPointInTime

Restores a DB cluster to an arbitrary point in time. Users can restore to any point in time before `LatestRestorableTime` for up to `BackupRetentionPeriod` days. The target DB cluster is created from the source DB cluster with the same configuration as the original DB cluster, except that the new DB cluster is created with the default DB security group.

## Note

This action only restores the DB cluster, not the DB instances for that DB cluster. You must invoke the [CreateDBInstance](#) action to create DB instances for the restored DB cluster, specifying the identifier of the restored DB cluster in `DBClusterIdentifier`. You can create DB instances only after the `RestoreDBClusterToPointInTime` action has completed and the DB cluster is available.

## Request Parameters

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

### **DBClusterIdentifier**

The name of the new DB cluster to be created.

Constraints:

- Must contain from 1 to 63 letters, numbers, or hyphens
- First character must be a letter
- Cannot end with a hyphen or contain two consecutive hyphens

Type: String

Required: Yes

### **DBClusterParameterGroupName**

The name of the DB cluster parameter group to associate with the new DB cluster.

Constraints:

- If supplied, must match the name of an existing `DBClusterParameterGroup`.

Type: String

Required: No

### **DBSubnetGroupName**

The DB subnet group name to use for the new DB cluster.

Constraints: If supplied, must match the name of an existing DBSubnetGroup.

Example: mySubnetgroup

Type: String

Required: No

### **DeletionProtection**

A value that indicates whether the DB cluster has deletion protection enabled. The database can't be deleted when deletion protection is enabled. By default, deletion protection is disabled.

Type: Boolean

Required: No

### **EnableCloudwatchLogsExports.member.N**

The list of logs that the restored DB cluster is to export to CloudWatch Logs.

Type: Array of strings

Required: No

### **EnableIAMDatabaseAuthentication**

True to enable mapping of Amazon Identity and Access Management (IAM) accounts to database accounts, and otherwise false.

Default: false

Type: Boolean

Required: No

### **KmsKeyId**

The Amazon KMS key identifier to use when restoring an encrypted DB cluster from an encrypted DB cluster.

The KMS key identifier is the Amazon Resource Name (ARN) for the KMS encryption key. If you are restoring a DB cluster with the same Amazon account that owns the KMS encryption key used to encrypt the new DB cluster, then you can use the KMS key alias instead of the ARN for the KMS encryption key.

You can restore to a new DB cluster and encrypt the new DB cluster with a KMS key that is different than the KMS key used to encrypt the source DB cluster. The new DB cluster is encrypted with the KMS key identified by the `KmsKeyId` parameter.

If you do not specify a value for the `KmsKeyId` parameter, then the following will occur:

- If the DB cluster is encrypted, then the restored DB cluster is encrypted using the KMS key that was used to encrypt the source DB cluster.
- If the DB cluster is not encrypted, then the restored DB cluster is not encrypted.

If `DBClusterIdentifier` refers to a DB cluster that is not encrypted, then the restore request is rejected.

Type: String

Required: No

### **OptionGroupName**

*(Not supported by Neptune)*

Type: String

Required: No

### **Port**

The port number on which the new DB cluster accepts connections.

Constraints: Value must be 1150-65535

Default: The same port as the original DB cluster.

Type: Integer

Required: No

### **RestoreToTime**

The date and time to restore the DB cluster to.

**Valid Values:** Value must be a time in Universal Coordinated Time (UTC) format

**Constraints:**

- Must be before the latest restorable time for the DB instance
- Must be specified if `UseLatestRestorableTime` parameter is not provided
- Cannot be specified if `UseLatestRestorableTime` parameter is true
- Cannot be specified if `RestoreType` parameter is `copy-on-write`

**Example:** `2015-03-07T23:45:00Z`

**Type:** Timestamp

**Required:** No

### **RestoreType**

The type of restore to be performed. You can specify one of the following values:

- `full-copy` - The new DB cluster is restored as a full copy of the source DB cluster.
- `copy-on-write` - The new DB cluster is restored as a clone of the source DB cluster.

If you don't specify a `RestoreType` value, then the new DB cluster is restored as a full copy of the source DB cluster.

**Type:** String

**Required:** No

### **ServerlessV2ScalingConfiguration**

Contains the scaling configuration of a Neptune Serverless DB cluster.

For more information, see [Using Amazon Neptune Serverless](#) in the *Amazon Neptune User Guide*.

**Type:** [ServerlessV2ScalingConfiguration](#) object

**Required:** No

### **SourceDBClusterIdentifier**

The identifier of the source DB cluster from which to restore.

**Constraints:**

- Must match the identifier of an existing DBCluster.

Type: String

Required: Yes

**StorageType**

Specifies the storage type to be associated with the DB cluster.

Valid values: `standard`, `iopt1`

Default: `standard`

Type: String

Required: No

**Tags.Tag.N**

The tags to be applied to the restored DB cluster.

Type: Array of [Tag](#) objects

Required: No

**UseLatestRestorableTime**

A value that is set to `true` to restore the DB cluster to the latest restorable backup time, and `false` otherwise.

Default: `false`

Constraints: Cannot be specified if `RestoreToTime` parameter is provided.

Type: Boolean

Required: No

**VpcSecurityGroupIds.VpcSecurityGroupId.N**

A list of VPC security groups that the new DB cluster belongs to.

Type: Array of strings

Required: No

## Response Elements

The following element is returned by the service.

### **DBCluster**

Contains the details of an Amazon Neptune DB cluster.

This data type is used as a response element in the [DescribeDBClusters](#).

Type: [DBCluster](#) object

## Errors

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

### **DBClusterAlreadyExistsFault**

User already has a DB cluster with the given identifier.

HTTP Status Code: 400

### **DBClusterNotFoundFault**

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

### **DBClusterParameterGroupNotFound**

*DBClusterParameterGroupName* does not refer to an existing DB Cluster parameter group.

HTTP Status Code: 404

### **DBClusterQuotaExceededFault**

User attempted to create a new DB cluster and the user has already reached the maximum allowed DB cluster quota.

HTTP Status Code: 403

**DBClusterSnapshotNotFoundFault**

*DBClusterSnapshotIdentifier* does not refer to an existing DB cluster snapshot.

HTTP Status Code: 404

**DBSubnetGroupNotFoundFault**

*DBSubnetGroupName* does not refer to an existing DB subnet group.

HTTP Status Code: 404

**InsufficientDBClusterCapacityFault**

The DB cluster does not have enough capacity for the current operation.

HTTP Status Code: 403

**InsufficientStorageClusterCapacity**

There is insufficient storage available for the current action. You may be able to resolve this error by updating your subnet group to use different Availability Zones that have more storage available.

HTTP Status Code: 400

**InvalidDBClusterSnapshotStateFault**

The supplied value is not a valid DB cluster snapshot state.

HTTP Status Code: 400

**InvalidDBClusterStateFault**

The DB cluster is not in a valid state.

HTTP Status Code: 400

**InvalidDBSnapshotState**

The state of the DB snapshot does not allow deletion.

HTTP Status Code: 400

**InvalidRestoreFault**

Cannot restore from vpc backup to non-vpc DB instance.

HTTP Status Code: 400

### **InvalidSubnet**

The requested subnet is invalid, or multiple subnets were requested that are not all in a common VPC.

HTTP Status Code: 400

### **InvalidVPCNetworkStateFault**

DB subnet group does not cover all Availability Zones after it is created because users' change.

HTTP Status Code: 400

### **KMSKeyNotAccessibleFault**

Error accessing KMS key.

HTTP Status Code: 400

### **OptionGroupNotFoundFault**

The designated option group could not be found.

HTTP Status Code: 404

### **StorageQuotaExceeded**

Request would result in user exceeding the allowed amount of storage available across all DB instances.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# StartDBCluster

Starts an Amazon Neptune DB cluster that was stopped using the Amazon console, the Amazon CLI `stop-db-cluster` command, or the `StopDBCluster` API.

## Request Parameters

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

### DBClusterIdentifier

The DB cluster identifier of the Neptune DB cluster to be started. This parameter is stored as a lowercase string.

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### DBCluster

Contains the details of an Amazon Neptune DB cluster.

This data type is used as a response element in the [DescribeDBClusters](#).

Type: [DBCluster](#) object

## Errors

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

### DBClusterNotFoundFault

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

## InvalidDBClusterStateFault

The DB cluster is not in a valid state.

HTTP Status Code: 400

## InvalidDBInstanceState

The specified DB instance is not in the *available* state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# StopDBCluster

Stops an Amazon Neptune DB cluster. When you stop a DB cluster, Neptune retains the DB cluster's metadata, including its endpoints and DB parameter groups.

Neptune also retains the transaction logs so you can do a point-in-time restore if necessary.

## Request Parameters

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

### DBClusterIdentifier

The DB cluster identifier of the Neptune DB cluster to be stopped. This parameter is stored as a lowercase string.

Type: String

Required: Yes

## Response Elements

The following element is returned by the service.

### DBCluster

Contains the details of an Amazon Neptune DB cluster.

This data type is used as a response element in the [DescribeDBClusters](#).

Type: [DBCluster](#) object

## Errors

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

### DBClusterNotFoundFault

*DBClusterIdentifier* does not refer to an existing DB cluster.

HTTP Status Code: 404

## InvalidDBClusterStateFault

The DB cluster is not in a valid state.

HTTP Status Code: 400

## InvalidDBInstanceState

The specified DB instance is not in the *available* state.

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 PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# Data Types

The Amazon Neptune 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:

- [AvailabilityZone](#)
- [CharacterSet](#)
- [CloudwatchLogsExportConfiguration](#)
- [ClusterPendingModifiedValues](#)
- [DBCluster](#)
- [DBClusterEndpoint](#)
- [DBClusterMember](#)
- [DBClusterOptionGroupStatus](#)
- [DBClusterParameterGroup](#)
- [DBClusterRole](#)
- [DBClusterSnapshot](#)
- [DBClusterSnapshotAttribute](#)
- [DBClusterSnapshotAttributesResult](#)
- [DBEngineVersion](#)
- [DBInstance](#)
- [DBInstanceStatusInfo](#)
- [DBParameterGroup](#)
- [DBParameterGroupStatus](#)
- [DBSecurityGroupMembership](#)
- [DBSubnetGroup](#)

- [DomainMembership](#)
- [DoubleRange](#)
- [Endpoint](#)
- [EngineDefaults](#)
- [Event](#)
- [EventCategoriesMap](#)
- [EventSubscription](#)
- [Filter](#)
- [GlobalCluster](#)
- [GlobalClusterMember](#)
- [OptionGroupMembership](#)
- [OrderableDBInstanceOption](#)
- [Parameter](#)
- [PendingCloudwatchLogsExports](#)
- [PendingMaintenanceAction](#)
- [PendingModifiedValues](#)
- [Range](#)
- [ResourcePendingMaintenanceActions](#)
- [ServerlessV2ScalingConfiguration](#)
- [ServerlessV2ScalingConfigurationInfo](#)
- [Subnet](#)
- [Tag](#)
- [Timezone](#)
- [UpgradeTarget](#)
- [ValidDBInstanceModificationsMessage](#)
- [ValidStorageOptions](#)
- [VpcSecurityGroupMembership](#)

# AvailabilityZone

Specifies an Availability Zone.

## Contents

### 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](#)



# CharacterSet

Specifies a character set.

## Contents

### CharacterSetDescription

The description of the character set.

Type: String

Required: No

### CharacterSetName

The name of the character set.

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](#)

# CloudwatchLogsExportConfiguration

The configuration setting for the log types to be enabled for export to CloudWatch Logs for a specific DB instance or DB cluster.

The `EnableLogTypes` and `DisableLogTypes` arrays determine which logs will be exported (or not exported) to CloudWatch Logs.

Valid log types are: `audit` (to publish audit logs) and `slowquery` (to publish slow-query logs). See [Publishing Neptune logs to Amazon CloudWatch logs](#).

## Contents

### `DisableLogTypes.member.N`

The list of log types to disable.

Type: Array of strings

Required: No

### `EnableLogTypes.member.N`

The list of log types to enable.

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](#)

# ClusterPendingModifiedValues

This data type is used as a response element in the `ModifyDBCluster` operation and contains changes that will be applied during the next maintenance window.

## Contents

### AllocatedStorage

The allocated storage size in gibibytes (GiB) for database engines. For Neptune, `AllocatedStorage` always returns 1, because Neptune DB cluster storage size isn't fixed, but instead automatically adjusts as needed.

Type: Integer

Required: No

### BackupRetentionPeriod

The number of days for which automatic DB snapshots are retained.

Type: Integer

Required: No

### DBClusterIdentifier

The `DBClusterIdentifier` value for the DB cluster.

Type: String

Required: No

### EngineVersion

The database engine version.

Type: String

Required: No

### IAMDatabaseAuthenticationEnabled

A value that indicates whether mapping of AWS Identity and Access Management (IAM) accounts to database accounts is enabled.

Type: Boolean

Required: No

## lops

The Provisioned IOPS (I/O operations per second) value. This setting is only for Multi-AZ DB clusters.

Type: Integer

Required: No

## PendingCloudwatchLogsExports

This PendingCloudwatchLogsExports structure specifies pending changes to which CloudWatch logs are enabled and which are disabled.

Type: [PendingCloudwatchLogsExports](#) object

Required: No

## StorageType

The pending change in storage type for the DB cluster. Valid Values:

- **standard** – ( *the default* ) Configures cost-effective database storage for applications with moderate to small I/O usage.
- **iopt1** – Enables [I/O-Optimized storage](#) that's designed to meet the needs of I/O-intensive graph workloads that require predictable pricing with low I/O latency and consistent I/O throughput.

Neptune I/O-Optimized storage is only available starting with engine release 1.3.0.0.

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](#)

# DBCluster

Contains the details of an Amazon Neptune DB cluster.

This data type is used as a response element in the [DescribeDBClusters](#).

## Contents

### AllocatedStorage

`AllocatedStorage` always returns 1, because Neptune DB cluster storage size is not fixed, but instead automatically adjusts as needed.

Type: Integer

Required: No

### AssociatedRoles.DBClusterRole.N

Provides a list of the Amazon Identity and Access Management (IAM) roles that are associated with the DB cluster. IAM roles that are associated with a DB cluster grant permission for the DB cluster to access other Amazon services on your behalf.

Type: Array of [DBClusterRole](#) objects

Required: No

### AutomaticRestartTime

Time at which the DB cluster will be automatically restarted.

Type: Timestamp

Required: No

### AvailabilityZones.AvailabilityZone.N

Provides the list of EC2 Availability Zones that instances in the DB cluster can be created in.

Type: Array of strings

Required: No

## **BackupRetentionPeriod**

Specifies the number of days for which automatic DB snapshots are retained.

Type: Integer

Required: No

## **CharacterSetName**

Not supported by Neptune.

Type: String

Required: No

## **CloneGroupId**

Identifies the clone group to which the DB cluster is associated.

Type: String

Required: No

## **ClusterCreateTime**

Specifies the time when the DB cluster was created, in Universal Coordinated Time (UTC).

Type: Timestamp

Required: No

## **CopyTagsToSnapshot**

*If set to `true`, tags are copied to any snapshot of the DB cluster that is created.*

Type: Boolean

Required: No

## **CrossAccountClone**

If set to `true`, the DB cluster can be cloned across accounts.

Type: Boolean

Required: No

**DatabaseName**

Contains the name of the initial database of this DB cluster that was provided at create time, if one was specified when the DB cluster was created. This same name is returned for the life of the DB cluster.

Type: String

Required: No

**DBClusterArn**

The Amazon Resource Name (ARN) for the DB cluster.

Type: String

Required: No

**DBClusterIdentifier**

Contains a user-supplied DB cluster identifier. This identifier is the unique key that identifies a DB cluster.

Type: String

Required: No

**DBClusterMembers.DBClusterMember.N**

Provides the list of instances that make up the DB cluster.

Type: Array of [DBClusterMember](#) objects

Required: No

**DBClusterOptionGroupMemberships.DBClusterOptionGroup.N**

Not supported by Neptune.

Type: Array of [DBClusterOptionGroupStatus](#) objects

Required: No

**DBClusterParameterGroup**

Specifies the name of the DB cluster parameter group for the DB cluster.



Type: String

Required: No

### **DbClusterResourceId**

The Amazon Region-unique, immutable identifier for the DB cluster. This identifier is found in Amazon CloudTrail log entries whenever the Amazon KMS key for the DB cluster is accessed.

Type: String

Required: No

### **DBSubnetGroup**

Specifies information on the subnet group associated with the DB cluster, including the name, description, and subnets in the subnet group.

Type: String

Required: No

### **DeletionProtection**

Indicates whether or not the DB cluster has deletion protection enabled. The database can't be deleted when deletion protection is enabled.

Type: Boolean

Required: No

### **EarliestRestorableTime**

Specifies the earliest time to which a database can be restored with point-in-time restore.

Type: Timestamp

Required: No

### **EnabledCloudwatchLogsExports.member.N**

A list of the log types that this DB cluster is configured to export to CloudWatch Logs. Valid log types are: `audit` (to publish audit logs to CloudWatch) and `slowquery` (to publish slow-query logs to CloudWatch). See [Publishing Neptune logs to Amazon CloudWatch logs](#).

Type: Array of strings

Required: No

### **Endpoint**

Specifies the connection endpoint for the primary instance of the DB cluster.

Type: String

Required: No

### **Engine**

Provides the name of the database engine to be used for this DB cluster.

Type: String

Required: No

### **EngineVersion**

Indicates the database engine version.

Type: String

Required: No

### **GlobalClusterIdentifier**

Contains a user-supplied global database cluster identifier. This identifier is the unique key that identifies a global database.

Type: String

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

Pattern: `[A-Za-z][0-9A-Za-z-:._]*`

Required: No

### **HostedZoneId**

Specifies the ID that Amazon Route 53 assigns when you create a hosted zone.

Type: String

Required: No

## **IAMDatabaseAuthenticationEnabled**

True if mapping of Amazon Identity and Access Management (IAM) accounts to database accounts is enabled, and otherwise false.

Type: Boolean

Required: No

## **IOOptimizedNextAllowedModificationTime**

The next time you can modify the DB cluster to use the `iopt1` storage type.

Type: Timestamp

Required: No

## **KmsKeyId**

If `StorageEncrypted` is true, the Amazon KMS key identifier for the encrypted DB cluster.

Type: String

Required: No

## **LatestRestorableTime**

Specifies the latest time to which a database can be restored with point-in-time restore.

Type: Timestamp

Required: No

## **MasterUsername**

Not supported by Neptune.

Type: String

Required: No

## **MultiAZ**

Specifies whether the DB cluster has instances in multiple Availability Zones.

Type: Boolean

Required: No

## PendingModifiedValues

This data type is used as a response element in the `ModifyDBCluster` operation and contains changes that will be applied during the next maintenance window.

Type: [ClusterPendingModifiedValues](#) object

Required: No

## PercentProgress

Specifies the progress of the operation as a percentage.

Type: String

Required: No

## Port

Specifies the port that the database engine is listening on.

Type: Integer

Required: No

## PreferredBackupWindow

Specifies the daily time range during which automated backups are created if automated backups are enabled, as determined by the `BackupRetentionPeriod`.

Type: String

Required: No

## PreferredMaintenanceWindow

Specifies the weekly time range during which system maintenance can occur, in Universal Coordinated Time (UTC).

Type: String

Required: No

## ReaderEndpoint

The reader endpoint for the DB cluster. The reader endpoint for a DB cluster load-balances connections across the Read Replicas that are available in a DB cluster. As clients request new

connections to the reader endpoint, Neptune distributes the connection requests among the Read Replicas in the DB cluster. This functionality can help balance your read workload across multiple Read Replicas in your DB cluster.

If a failover occurs, and the Read Replica that you are connected to is promoted to be the primary instance, your connection is dropped. To continue sending your read workload to other Read Replicas in the cluster, you can then reconnect to the reader endpoint.

Type: String

Required: No

### **ReadReplicaIdentifiers.ReadReplicaIdentifier.N**

Contains one or more identifiers of the Read Replicas associated with this DB cluster.

Type: Array of strings

Required: No

### **ReplicationSourceIdentifier**

Not supported by Neptune.

Type: String

Required: No

### **ServerlessV2ScalingConfiguration**

Shows the scaling configuration for a Neptune Serverless DB cluster.

For more information, see [Using Amazon Neptune Serverless](#) in the *Amazon Neptune User Guide*.

Type: [ServerlessV2ScalingConfigurationInfo](#) object

Required: No

### **Status**

Specifies the current state of this DB cluster.

Type: String

Required: No

## StorageEncrypted

Specifies whether the DB cluster is encrypted.

Type: Boolean

Required: No

## StorageType

The storage type used by the DB cluster.

Valid Values:

- **standard** – ( *the default* ) Provides cost-effective database storage for applications with moderate to small I/O usage.
- **iopt1** – Enables [I/O-Optimized storage](#) that's designed to meet the needs of I/O-intensive graph workloads that require predictable pricing with low I/O latency and consistent I/O throughput.

Neptune I/O-Optimized storage is only available starting with engine release 1.3.0.0.

Type: String

Required: No

## VpcSecurityGroups.VpcSecurityGroupMembership.N

Provides a list of VPC security groups that the DB cluster belongs to.

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



# DBClusterEndpoint

This data type represents the information you need to connect to an Amazon Neptune DB cluster. This data type is used as a response element in the following actions:

- `CreateDBClusterEndpoint`
- `DescribeDBClusterEndpoints`
- `ModifyDBClusterEndpoint`
- `DeleteDBClusterEndpoint`

For the data structure that represents Amazon Neptune DB instance endpoints, see `Endpoint`.

## Contents

### CustomEndpointType

The type associated with a custom endpoint. One of: `READER`, `WRITER`, `ANY`.

Type: String

Required: No

### DBClusterEndpointArn

The Amazon Resource Name (ARN) for the endpoint.

Type: String

Required: No

### DBClusterEndpointIdentifier

The identifier associated with the endpoint. This parameter is stored as a lowercase string.

Type: String

Required: No

### DBClusterEndpointResourceIdentifier

A unique system-generated identifier for an endpoint. It remains the same for the whole life of the endpoint.



Type: String

Required: No

### **DBClusterIdentifier**

The DB cluster identifier of the DB cluster associated with the endpoint. This parameter is stored as a lowercase string.

Type: String

Required: No

### **Endpoint**

The DNS address of the endpoint.

Type: String

Required: No

### **EndpointType**

The type of the endpoint. One of: READER, WRITER, CUSTOM.

Type: String

Required: No

### **ExcludedMembers.member.N**

List of DB instance identifiers that aren't part of the custom endpoint group. All other eligible instances are reachable through the custom endpoint. Only relevant if the list of static members is empty.

Type: Array of strings

Required: No

### **StaticMembers.member.N**

List of DB instance identifiers that are part of the custom endpoint group.

Type: Array of strings

Required: No

## Status

The current status of the endpoint. One of: `creating`, `available`, `deleting`, `inactive`, `modifying`. The `inactive` state applies to an endpoint that cannot be used for a certain kind of cluster, such as a `writer` endpoint for a read-only secondary cluster in a global database.

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](#)

# DBClusterMember

Contains information about an instance that is part of a DB cluster.

## Contents

### DBClusterParameterGroupStatus

Specifies the status of the DB cluster parameter group for this member of the DB cluster.

Type: String

Required: No

### DBInstanceIdentifier

Specifies the instance identifier for this member of the DB cluster.

Type: String

Required: No

### IsClusterWriter

Value that is `true` if the cluster member is the primary instance for the DB cluster and `false` otherwise.

Type: Boolean

Required: No

### PromotionTier

A value that specifies the order in which a Read Replica is promoted to the primary instance after a failure of the existing primary instance.

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](#)

# DBClusterOptionGroupStatus

Not supported by Neptune.

## Contents

### DBClusterOptionGroupName

Not supported by Neptune.

Type: String

Required: No

### Status

Not supported by Neptune.

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](#)

# DBClusterParameterGroup

Contains the details of an Amazon Neptune DB cluster parameter group.

This data type is used as a response element in the [DescribeDBClusterParameterGroups](#) action.

## Contents

### DBClusterParameterGroupArn

The Amazon Resource Name (ARN) for the DB cluster parameter group.

Type: String

Required: No

### DBClusterParameterGroupName

Provides the name of the DB cluster parameter group.

Type: String

Required: No

### DBParameterGroupFamily

Provides the name of the DB parameter group family that this DB cluster parameter group is compatible with.

Type: String

Required: No

### Description

Provides the customer-specified description for this DB cluster parameter 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](#)

# DBClusterRole

Describes an Amazon Identity and Access Management (IAM) role that is associated with a DB cluster.

## Contents

### FeatureName

The name of the feature associated with the Amazon Identity and Access Management (IAM) role. For the list of supported feature names, see [DescribeDBEngineVersions](#).

Type: String

Required: No

### RoleArn

The Amazon Resource Name (ARN) of the IAM role that is associated with the DB cluster.

Type: String

Required: No

### Status

Describes the state of association between the IAM role and the DB cluster. The Status property returns one of the following values:

- **ACTIVE** - the IAM role ARN is associated with the DB cluster and can be used to access other Amazon services on your behalf.
- **PENDING** - the IAM role ARN is being associated with the DB cluster.
- **INVALID** - the IAM role ARN is associated with the DB cluster, but the DB cluster is unable to assume the IAM role in order to access other Amazon services on your behalf.

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](#)

# DBClusterSnapshot

Contains the details for an Amazon Neptune DB cluster snapshot.

This data type is used as a response element in the [DescribeDBClusterSnapshots](#) action.

## Contents

### AllocatedStorage

Specifies the allocated storage size in gibibytes (GiB).

Type: Integer

Required: No

### AvailabilityZones.AvailabilityZone.N

Provides the list of EC2 Availability Zones that instances in the DB cluster snapshot can be restored in.

Type: Array of strings

Required: No

### ClusterCreateTime

Specifies the time when the DB cluster was created, in Universal Coordinated Time (UTC).

Type: Timestamp

Required: No

### DBClusterIdentifier

Specifies the DB cluster identifier of the DB cluster that this DB cluster snapshot was created from.

Type: String

Required: No

### DBClusterSnapshotArn

The Amazon Resource Name (ARN) for the DB cluster snapshot.

Type: String

Required: No

### **DBClusterSnapshotIdentifier**

Specifies the identifier for a DB cluster snapshot. Must match the identifier of an existing snapshot.

After you restore a DB cluster using a `DBClusterSnapshotIdentifier`, you must specify the same `DBClusterSnapshotIdentifier` for any future updates to the DB cluster. When you specify this property for an update, the DB cluster is not restored from the snapshot again, and the data in the database is not changed.

However, if you don't specify the `DBClusterSnapshotIdentifier`, an empty DB cluster is created, and the original DB cluster is deleted. If you specify a property that is different from the previous snapshot restore property, the DB cluster is restored from the snapshot specified by the `DBClusterSnapshotIdentifier`, and the original DB cluster is deleted.

Type: String

Required: No

### **Engine**

Specifies the name of the database engine.

Type: String

Required: No

### **EngineVersion**

Provides the version of the database engine for this DB cluster snapshot.

Type: String

Required: No

### **IAMDatabaseAuthenticationEnabled**

True if mapping of Amazon Identity and Access Management (IAM) accounts to database accounts is enabled, and otherwise false.

Type: Boolean

Required: No

### **KmsKeyId**

If `StorageEncrypted` is true, the Amazon KMS key identifier for the encrypted DB cluster snapshot.

Type: String

Required: No

### **LicenseModel**

Provides the license model information for this DB cluster snapshot.

Type: String

Required: No

### **MasterUsername**

Not supported by Neptune.

Type: String

Required: No

### **PercentProgress**

Specifies the percentage of the estimated data that has been transferred.

Type: Integer

Required: No

### **Port**

Specifies the port that the DB cluster was listening on at the time of the snapshot.

Type: Integer

Required: No

### **SnapshotCreateTime**

Provides the time when the snapshot was taken, in Universal Coordinated Time (UTC).

Type: Timestamp

Required: No

### **SnapshotType**

Provides the type of the DB cluster snapshot.

Type: String

Required: No

### **SourceDBClusterSnapshotArn**

If the DB cluster snapshot was copied from a source DB cluster snapshot, the Amazon Resource Name (ARN) for the source DB cluster snapshot, otherwise, a null value.

Type: String

Required: No

### **Status**

Specifies the status of this DB cluster snapshot.

Type: String

Required: No

### **StorageEncrypted**

Specifies whether the DB cluster snapshot is encrypted.

Type: Boolean

Required: No

### **StorageType**

The storage type associated with the DB cluster snapshot.

Type: String

Required: No

### **VpcId**

Provides the VPC ID associated with the DB cluster snapshot.

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](#)

# DBClusterSnapshotAttribute

Contains the name and values of a manual DB cluster snapshot attribute.

Manual DB cluster snapshot attributes are used to authorize other Amazon accounts to restore a manual DB cluster snapshot. For more information, see the [ModifyDBClusterSnapshotAttribute](#) API action.

## Contents

### AttributeName

The name of the manual DB cluster snapshot attribute.

The attribute named `restore` refers to the list of Amazon accounts that have permission to copy or restore the manual DB cluster snapshot. For more information, see the [ModifyDBClusterSnapshotAttribute](#) API action.

Type: String

Required: No

### AttributeValues.AttributeValue.N

The value(s) for the manual DB cluster snapshot attribute.

If the `AttributeName` field is set to `restore`, then this element returns a list of IDs of the Amazon accounts that are authorized to copy or restore the manual DB cluster snapshot. If a value of `all` is in the list, then the manual DB cluster snapshot is public and available for any Amazon account to copy or restore.

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](#)



# DBClusterSnapshotAttributesResult

Contains the results of a successful call to the [DescribeDBClusterSnapshotAttributes](#) API action.

Manual DB cluster snapshot attributes are used to authorize other Amazon accounts to copy or restore a manual DB cluster snapshot. For more information, see the [ModifyDBClusterSnapshotAttribute](#) API action.

## Contents

### DBClusterSnapshotAttributes.DBClusterSnapshotAttribute.N

The list of attributes and values for the manual DB cluster snapshot.

Type: Array of [DBClusterSnapshotAttribute](#) objects

Required: No

### DBClusterSnapshotIdentifier

The identifier of the manual DB cluster snapshot that the attributes apply to.

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](#)

# DBEngineVersion

This data type is used as a response element in the action [DescribeDBEngineVersions](#).

## Contents

### DBEngineDescription

The description of the database engine.

Type: String

Required: No

### DBEngineVersionDescription

The description of the database engine version.

Type: String

Required: No

### DBParameterGroupFamily

The name of the DB parameter group family for the database engine.

Type: String

Required: No

### DefaultCharacterSet

*(Not supported by Neptune)*

Type: [CharacterSet](#) object

Required: No

### Engine

The name of the database engine.

Type: String

Required: No

## EngineVersion

The version number of the database engine.

Type: String

Required: No

## ExportableLogTypes.member.N

The types of logs that the database engine has available for export to CloudWatch Logs.

Type: Array of strings

Required: No

## SupportedCharacterSets.CharacterSet.N

*(Not supported by Neptune)*

Type: Array of [CharacterSet](#) objects

Required: No

## SupportedTimezones.Timezone.N

A list of the time zones supported by this engine for the Timezone parameter of the CreateDBInstance action.

Type: Array of [Timezone](#) objects

Required: No

## SupportsGlobalDatabases

A value that indicates whether you can use Aurora global databases with a specific DB engine version.

Type: Boolean

Required: No

## SupportsLogExportsToCloudwatchLogs

A value that indicates whether the engine version supports exporting the log types specified by ExportableLogTypes to CloudWatch Logs.

Type: Boolean

Required: No

### **SupportsReadReplica**

Indicates whether the database engine version supports read replicas.

Type: Boolean

Required: No

### **ValidUpgradeTarget.UpgradeTarget.N**

A list of engine versions that this database engine version can be upgraded to.

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

# DBInstance

Contains the details of an Amazon Neptune DB instance.

This data type is used as a response element in the [DescribeDBInstances](#) action.

## Contents

### AllocatedStorage

Not supported by Neptune.

Type: Integer

Required: No

### AutoMinorVersionUpgrade

Indicates that minor version patches are applied automatically.

Type: Boolean

Required: No

### AvailabilityZone

Specifies the name of the Availability Zone the DB instance is located in.

Type: String

Required: No

### BackupRetentionPeriod

Specifies the number of days for which automatic DB snapshots are retained.

Type: Integer

Required: No

### CACertificateIdentifier

The identifier of the CA certificate for this DB instance.

Type: String

Required: No

### **CharacterSetName**

*(Not supported by Neptune)*

Type: String

Required: No

### **CopyTagsToSnapshot**

Specifies whether tags are copied from the DB instance to snapshots of the DB instance.

Type: Boolean

Required: No

### **DBClusterIdentifier**

If the DB instance is a member of a DB cluster, contains the name of the DB cluster that the DB instance is a member of.

Type: String

Required: No

### **DBInstanceArn**

The Amazon Resource Name (ARN) for the DB instance.

Type: String

Required: No

### **DBInstanceClass**

Contains the name of the compute and memory capacity class of the DB instance.

Type: String

Required: No

### **DBInstanceIdentifier**

Contains a user-supplied database identifier. This identifier is the unique key that identifies a DB instance.

Type: String

Required: No

### **DbInstancePort**

Specifies the port that the DB instance listens on. If the DB instance is part of a DB cluster, this can be a different port than the DB cluster port.

Type: Integer

Required: No

### **DBInstanceStatus**

Specifies the current state of this database.

Type: String

Required: No

### **DbiResourceId**

The Amazon Region-unique, immutable identifier for the DB instance. This identifier is found in Amazon CloudTrail log entries whenever the Amazon KMS key for the DB instance is accessed.

Type: String

Required: No

### **DBName**

The database name.

Type: String

Required: No

### **DBParameterGroups.DBParameterGroup.N**

Provides the list of DB parameter groups applied to this DB instance.

Type: Array of [DBParameterGroupStatus](#) objects

Required: No

## **DBSecurityGroups.DBSecurityGroup.N**

Provides List of DB security group elements containing only `DBSecurityGroup.Name` and `DBSecurityGroup.Status` subelements.

Type: Array of [DBSecurityGroupMembership](#) objects

Required: No

## **DBSubnetGroup**

Specifies information on the subnet group associated with the DB instance, including the name, description, and subnets in the subnet group.

Type: [DBSubnetGroup](#) object

Required: No

## **DeletionProtection**

Indicates whether or not the DB instance has deletion protection enabled. The instance can't be deleted when deletion protection is enabled. See [Deleting a DB Instance](#).

Type: Boolean

Required: No

## **DomainMemberships.DomainMembership.N**

Not supported

Type: Array of [DomainMembership](#) objects

Required: No

## **EnabledCloudwatchLogsExports.member.N**

A list of log types that this DB instance is configured to export to CloudWatch Logs.

Type: Array of strings

Required: No

## **Endpoint**

Specifies the connection endpoint.



Type: [Endpoint](#) object

Required: No

### **Engine**

Provides the name of the database engine to be used for this DB instance.

Type: String

Required: No

### **EngineVersion**

Indicates the database engine version.

Type: String

Required: No

### **EnhancedMonitoringResourceArn**

The Amazon Resource Name (ARN) of the Amazon CloudWatch Logs log stream that receives the Enhanced Monitoring metrics data for the DB instance.

Type: String

Required: No

### **IAMDatabaseAuthenticationEnabled**

True if Amazon Identity and Access Management (IAM) authentication is enabled, and otherwise false.

Type: Boolean

Required: No

### **InstanceCreateTime**

Provides the date and time the DB instance was created.

Type: Timestamp

Required: No

**Iops**

Specifies the Provisioned IOPS (I/O operations per second) value.

Type: Integer

Required: No

**KmsKeyId**

Not supported: The encryption for DB instances is managed by the DB cluster.

Type: String

Required: No

**LatestRestorableTime**

Specifies the latest time to which a database can be restored with point-in-time restore.

Type: Timestamp

Required: No

**LicenseModel**

License model information for this DB instance.

Type: String

Required: No

**MasterUsername**

Not supported by Neptune.

Type: String

Required: No

**MonitoringInterval**

The interval, in seconds, between points when Enhanced Monitoring metrics are collected for the DB instance.

Type: Integer

Required: No

### **MonitoringRoleArn**

The ARN for the IAM role that permits Neptune to send Enhanced Monitoring metrics to Amazon CloudWatch Logs.

Type: String

Required: No

### **MultiAZ**

Specifies if the DB instance is a Multi-AZ deployment.

Type: Boolean

Required: No

### **OptionGroupMemberships.OptionGroupMembership.N**

*(Not supported by Neptune)*

Type: Array of [OptionGroupMembership](#) objects

Required: No

### **PendingModifiedValues**

Specifies that changes to the DB instance are pending. This element is only included when changes are pending. Specific changes are identified by subelements.

Type: [PendingModifiedValues](#) object

Required: No

### **PerformanceInsightsEnabled**

*(Not supported by Neptune)*

Type: Boolean

Required: No

### **PerformanceInsightsKMSKeyId**

*(Not supported by Neptune)*

Type: String

Required: No

### **PreferredBackupWindow**

Specifies the daily time range during which automated backups are created if automated backups are enabled, as determined by the `BackupRetentionPeriod`.

Type: String

Required: No

### **PreferredMaintenanceWindow**

Specifies the weekly time range during which system maintenance can occur, in Universal Coordinated Time (UTC).

Type: String

Required: No

### **PromotionTier**

A value that specifies the order in which a Read Replica is promoted to the primary instance after a failure of the existing primary instance.

Type: Integer

Required: No

### **PubliclyAccessible**

*This member has been deprecated.*

This flag should no longer be used.

Type: Boolean

Required: No

### **ReadReplicaDBClusterIdentifiers.ReadReplicaDBClusterIdentifier.N**

Contains one or more identifiers of DB clusters that are Read Replicas of this DB instance.

Type: Array of strings

Required: No

### **ReadReplicaDBInstanceIdentifiers.ReadReplicaDBInstanceIdentifier.N**

Contains one or more identifiers of the Read Replicas associated with this DB instance.

Type: Array of strings

Required: No

### **ReadReplicaSourceDBInstanceIdentifier**

Contains the identifier of the source DB instance if this DB instance is a Read Replica.

Type: String

Required: No

### **SecondaryAvailabilityZone**

If present, specifies the name of the secondary Availability Zone for a DB instance with multi-AZ support.

Type: String

Required: No

### **StatusInfos.DBInstanceStatusInfo.N**

The status of a Read Replica. If the instance is not a Read Replica, this is blank.

Type: Array of [DBInstanceStatusInfo](#) objects

Required: No

### **StorageEncrypted**

Not supported: The encryption for DB instances is managed by the DB cluster.

Type: Boolean

Required: No

### **StorageType**

Specifies the storage type associated with the DB instance.

Type: String

Required: No

### **TdeCredentialArn**

The ARN from the key store with which the instance is associated for TDE encryption.

Type: String

Required: No

### **Timezone**

Not supported.

Type: String

Required: No

### **VpcSecurityGroups.VpcSecurityGroupMembership.N**

Provides a list of VPC security group elements that the DB instance belongs to.

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

# DBInstanceStatusInfo

Provides a list of status information for a DB instance.

## Contents

### Message

Details of the error if there is an error for the instance. If the instance is not in an error state, this value is blank.

Type: String

Required: No

### Normal

Boolean value that is true if the instance is operating normally, or false if the instance is in an error state.

Type: Boolean

Required: No

### Status

Status of the DB instance. For a StatusType of read replica, the values can be replicating, error, stopped, or terminated.

Type: String

Required: No

### StatusType

This value is currently "read replication."

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](#)



# DBParameterGroup

Contains the details of an Amazon Neptune DB parameter group.

This data type is used as a response element in the [DescribeDBParameterGroups](#) action.

## Contents

### DBParameterGroupArn

The Amazon Resource Name (ARN) for the DB parameter group.

Type: String

Required: No

### DBParameterGroupFamily

Provides the name of the DB parameter group family that this DB parameter group is compatible with.

Type: String

Required: No

### DBParameterGroupName

Provides the name of the DB parameter group.

Type: String

Required: No

### Description

Provides the customer-specified description for this DB parameter 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](#)

# DBParameterGroupStatus

The status of the DB parameter group.

This data type is used as a response element in the following actions:

- [CreateDBInstance](#)
- [DeleteDBInstance](#)
- [ModifyDBInstance](#)
- [RebootDBInstance](#)

## Contents

### DBParameterGroupName

The name of the DP 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](#)

# DBSecurityGroupMembership

Specifies membership in a designated DB security group.

## Contents

### DBSecurityGroupName

The name of the DB security group.

Type: String

Required: No

### Status

The status of the DB 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](#)

# DBSubnetGroup

Contains the details of an Amazon Neptune DB subnet group.

This data type is used as a response element in the [DescribeDBSubnetGroups](#) action.

## Contents

### DBSubnetGroupArn

The Amazon Resource Name (ARN) for the DB subnet group.

Type: String

Required: No

### DBSubnetGroupDescription

Provides the description of the DB subnet group.

Type: String

Required: No

### DBSubnetGroupName

The name of the DB subnet group.

Type: String

Required: No

### SubnetGroupStatus

Provides the status of the DB subnet group.

Type: String

Required: No

### Subnets.Subnet.N

Contains a list of [Subnet](#) elements.

Type: Array of [Subnet](#) objects

Required: No

## VpcId

Provides the VpcId of the DB 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](#)

# DomainMembership

An Active Directory Domain membership record associated with a DB instance.

## Contents

### Domain

The identifier of the Active Directory Domain.

Type: String

Required: No

### FQDN

The fully qualified domain name of the Active Directory Domain.

Type: String

Required: No

### IAMRoleName

The name of the IAM role to be used when making API calls to the Directory Service.

Type: String

Required: No

### Status

The status of the DB instance's Active Directory Domain membership, such as joined, pending-join, failed etc).

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](#)



# DoubleRange

A range of double values.

## Contents

### From

The minimum value in the range.

Type: Double

Required: No

### To

The maximum value in the range.

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](#)

# Endpoint

Specifies a connection endpoint.

For the data structure that represents Amazon Neptune DB cluster endpoints, see `DBClusterEndpoint`.

## Contents

### Address

Specifies the DNS address of the DB instance.

Type: String

Required: No

### HostedZoneId

Specifies the ID that Amazon Route 53 assigns when you create a hosted zone.

Type: String

Required: No

### Port

Specifies the port that the database 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

Contains the result of a successful invocation of the [DescribeEngineDefaultParameters](#) action.

## Contents

### DBParameterGroupFamily

Specifies the name of the DB parameter group family that the engine default parameters apply to.

Type: String

Required: No

### Marker

An optional pagination token provided by a previous EngineDefaults request. If this parameter is specified, the response includes only records beyond the marker, up to the value specified by `MaxRecords`.

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

This data type is used as a response element in the [DescribeEvents](#) action.

## Contents

### Date

Specifies the date and time of the event.

Type: Timestamp

Required: No

### EventCategories.EventCategory.N

Specifies the category for the event.

Type: Array of strings

Required: No

### Message

Provides the text of this event.

Type: String

Required: No

### SourceArn

The Amazon Resource Name (ARN) for the event.

Type: String

Required: No

### SourceIdentifier

Provides the identifier for the source of the event.

Type: String

Required: No

## SourceType

Specifies the source type for this event.

Type: String

Valid Values: `db-instance` | `db-parameter-group` | `db-security-group` | `db-snapshot` | `db-cluster` | `db-cluster-snapshot`

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](#)

# EventCategoriesMap

Contains the results of a successful invocation of the [DescribeEventCategories](#) action.

## Contents

### EventCategories.EventCategory.N

The event categories for the specified source type

Type: Array of strings

Required: No

### SourceType

The source type that the returned categories belong to

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](#)



# EventSubscription

Contains the results of a successful invocation of the [DescribeEventSubscriptions](#) action.

## Contents

### CustomerAwsId

The Amazon customer account associated with the event notification subscription.

Type: String

Required: No

### CustSubscriptionId

The event notification subscription Id.

Type: String

Required: No

### Enabled

A Boolean value indicating if the subscription is enabled. True indicates the subscription is enabled.

Type: Boolean

Required: No

### EventCategoriesList.EventCategory.N

A list of event categories for the event notification subscription.

Type: Array of strings

Required: No

### EventSubscriptionArn

The Amazon Resource Name (ARN) for the event subscription.

Type: String

Required: No

## **SnsTopicArn**

The topic ARN of the event notification subscription.

Type: String

Required: No

## **SourceIdsList.SourceId.N**

A list of source IDs for the event notification subscription.

Type: Array of strings

Required: No

## **SourceType**

The source type for the event notification subscription.

Type: String

Required: No

## **Status**

The status of the event notification subscription.

Constraints:

Can be one of the following: creating | modifying | deleting | active | no-permission | topic-not-exist

The status "no-permission" indicates that Neptune no longer has permission to post to the SNS topic. The status "topic-not-exist" indicates that the topic was deleted after the subscription was created.

Type: String

Required: No

## **SubscriptionCreationTime**

The time the event notification subscription was created.

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](#)

# Filter

This type is not currently supported.

## Contents

### Name

This parameter is not currently supported.

Type: String

Required: Yes

### Values.Value.N

This parameter is not currently supported.

Type: Array of strings

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](#)

# GlobalCluster

Contains the details of an Amazon Neptune global database.

This data type is used as a response element for the [CreateGlobalCluster](#), [DescribeGlobalClusters](#), [ModifyGlobalCluster](#), [DeleteGlobalCluster](#), [FailoverGlobalCluster](#), and [RemoveFromGlobalCluster](#) actions.

## Contents

### DeletionProtection

The deletion protection setting for the global database.

Type: Boolean

Required: No

### Engine

The Neptune database engine used by the global database ("neptune").

Type: String

Required: No

### EngineVersion

The Neptune engine version used by the global database.

Type: String

Required: No

### GlobalClusterArn

The Amazon Resource Name (ARN) for the global database.

Type: String

Required: No

### GlobalClusterIdentifier

Contains a user-supplied global database cluster identifier. This identifier is the unique key that identifies a global database.

Type: String

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

Pattern: `[A-Za-z][0-9A-Za-z- : ._-]*`

Required: No

### **GlobalClusterMembers.GlobalClusterMember.N**

A list of cluster ARNs and instance ARNs for all the DB clusters that are part of the global database.

Type: Array of [GlobalClusterMember](#) objects

Required: No

### **GlobalClusterResourceId**

An immutable identifier for the global database that is unique within in all regions. This identifier is found in CloudTrail log entries whenever the KMS key for the DB cluster is accessed.

Type: String

Required: No

### **Status**

Specifies the current state of this global database.

Type: String

Required: No

### **StorageEncrypted**

The storage encryption setting for the global database.

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](#)

# GlobalClusterMember

A data structure with information about any primary and secondary clusters associated with an Neptune global database.

## Contents

### DBClusterArn

The Amazon Resource Name (ARN) for each Neptune cluster.

Type: String

Required: No

### IsWriter

Specifies whether the Neptune cluster is the primary cluster (that is, has read-write capability) for the Neptune global database with which it is associated.

Type: Boolean

Required: No

### Readers.member.N

The Amazon Resource Name (ARN) for each read-only secondary cluster associated with the Neptune global database.

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](#)





# OptionGroupMembership

Not supported by Neptune.

## Contents

### OptionGroupName

Not supported by Neptune.

Type: String

Required: No

### Status

Not supported by Neptune.

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](#)

# OrderableDBInstanceOption

Contains a list of available options for a DB instance.

This data type is used as a response element in the [DescribeOrderableDBInstanceOptions](#) action.

## Contents

### AvailabilityZones.AvailabilityZone.N

A list of Availability Zones for a DB instance.

Type: Array of [AvailabilityZone](#) objects

Required: No

### DBInstanceClass

The DB instance class for a DB instance.

Type: String

Required: No

### Engine

The engine type of a DB instance.

Type: String

Required: No

### EngineVersion

The engine version of a DB instance.

Type: String

Required: No

### LicenseModel

The license model for a DB instance.

Type: String

Required: No

**MaxIopsPerDbInstance**

Maximum total provisioned IOPS for a DB instance.

Type: Integer

Required: No

**MaxIopsPerGib**

Maximum provisioned IOPS per GiB for a DB instance.

Type: Double

Required: No

**MaxStorageSize**

Maximum storage size for a DB instance.

Type: Integer

Required: No

**MinIopsPerDbInstance**

Minimum total provisioned IOPS for a DB instance.

Type: Integer

Required: No

**MinIopsPerGib**

Minimum provisioned IOPS per GiB for a DB instance.

Type: Double

Required: No

**MinStorageSize**

Minimum storage size for a DB instance.

Type: Integer

Required: No

### **MultiAZCapable**

Indicates whether a DB instance is Multi-AZ capable.

Type: Boolean

Required: No

### **ReadReplicaCapable**

Indicates whether a DB instance can have a Read Replica.

Type: Boolean

Required: No

### **StorageType**

Not applicable. In Neptune the storage type is managed at the DB Cluster level.

Type: String

Required: No

### **SupportsEnhancedMonitoring**

Indicates whether a DB instance supports Enhanced Monitoring at intervals from 1 to 60 seconds.

Type: Boolean

Required: No

### **SupportsGlobalDatabases**

A value that indicates whether you can use Neptune global databases with a specific combination of other DB engine attributes.

Type: Boolean

Required: No

### **SupportsIAMDatabaseAuthentication**

Indicates whether a DB instance supports IAM database authentication.

Type: Boolean

Required: No

### **SupportsIOPS**

Indicates whether a DB instance supports provisioned IOPS.

Type: Boolean

Required: No

### **SupportsPerformanceInsights**

*(Not supported by Neptune)*

Type: Boolean

Required: No

### **SupportsStorageEncryption**

Indicates whether a DB instance supports encrypted storage.

Type: Boolean

Required: No

### **Vpc**

Indicates whether a DB instance is in a VPC.

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](#)



# Parameter

Specifies a parameter.

## Contents

### AllowedValues

Specifies the valid range of values for the parameter.

Type: String

Required: No

### ApplyMethod

Indicates when to apply parameter updates.

Type: String

Valid Values: `immediate` | `pending-reboot`

Required: No

### ApplyType

Specifies the engine specific parameters type.

Type: String

Required: No

### DataType

Specifies the valid data type for the parameter.

Type: String

Required: No

### Description

Provides 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 engine version to which the parameter can apply.

Type: String

Required: No

## **ParameterName**

Specifies the name of the parameter.

Type: String

Required: No

## **ParameterValue**

Specifies the value of the parameter.

Type: String

Required: No

## **Source**

Indicates 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](#)

# PendingCloudwatchLogsExports

A list of the log types whose configuration is still pending. In other words, these log types are in the process of being activated or deactivated.

Valid log types are: `audit` (to publish audit logs) and `slowquery` (to publish slow-query logs). See [Publishing Neptune logs to Amazon CloudWatch logs](#).

## Contents

### **LogTypesToDisable.member.N**

Log types that are in the process of being enabled. After they are enabled, these log types are exported to CloudWatch Logs.

Type: Array of strings

Required: No

### **LogTypesToEnable.member.N**

Log types that are in the process of being deactivated. After they are deactivated, these log types aren't exported to CloudWatch Logs.

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](#)

# PendingMaintenanceAction

Provides information about a pending maintenance action for a resource.

## Contents

### Action

The type of pending maintenance action that is available for the resource.

Type: String

Required: No

### AutoAppliedAfterDate

The date of the maintenance window when the action is applied. The maintenance action is applied to the resource during its first maintenance window after this date. If this date is specified, any next-maintenance opt-in requests are ignored.

Type: Timestamp

Required: No

### CurrentApplyDate

The effective date when the pending maintenance action is applied to the resource. This date takes into account opt-in requests received from the [ApplyPendingMaintenanceAction](#) API, the `AutoAppliedAfterDate`, and the `ForcedApplyDate`. This value is blank if an opt-in request has not been received and nothing has been specified as `AutoAppliedAfterDate` or `ForcedApplyDate`.

Type: Timestamp

Required: No

### Description

A description providing more detail about the maintenance action.

Type: String

Required: No

## ForcedApplyDate

The date when the maintenance action is automatically applied. The maintenance action is applied to the resource on this date regardless of the maintenance window for the resource. If this date is specified, any `immediate` opt-in requests are ignored.

Type: Timestamp

Required: No

## OptInStatus

Indicates the type of opt-in request that has been received for the resource.

Type: String

Required: No

## See Also

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

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

# PendingModifiedValues

This data type is used as a response element in the [ModifyDBInstance](#) action.

## Contents

### AllocatedStorage

Contains the new `AllocatedStorage` size for the DB instance that will be applied or is currently being applied.

Type: Integer

Required: No

### BackupRetentionPeriod

Specifies the pending number of days for which automated backups are retained.

Type: Integer

Required: No

### CACertificateIdentifier

Specifies the identifier of the CA certificate for the DB instance.

Type: String

Required: No

### DBInstanceClass

Contains the new `DBInstanceClass` for the DB instance that will be applied or is currently being applied.

Type: String

Required: No

### DBInstanceIdentifier

Contains the new `DBInstanceIdentifier` for the DB instance that will be applied or is currently being applied.

Type: String

Required: No

### **DBSubnetGroupName**

The new DB subnet group for the DB instance.

Type: String

Required: No

### **EngineVersion**

Indicates the database engine version.

Type: String

Required: No

### **Iops**

Specifies the new Provisioned IOPS value for the DB instance that will be applied or is currently being applied.

Type: Integer

Required: No

### **LicenseModel**

Not supported by Neptune.

Type: String

Required: No

### **MasterUserPassword**

Not supported by Neptune.

Type: String

Required: No

### **MultiAZ**

Indicates that the Single-AZ DB instance is to change to a Multi-AZ deployment.

Type: Boolean

Required: No

### **PendingCloudwatchLogsExports**

This `PendingCloudwatchLogsExports` structure specifies pending changes to which CloudWatch logs are enabled and which are disabled.

Type: [PendingCloudwatchLogsExports](#) object

Required: No

### **Port**

Specifies the pending port for the DB instance.

Type: Integer

Required: No

### **StorageType**

Not applicable. In Neptune the storage type is managed at the DB Cluster level.

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](#)



# Range

A range of integer values.

## Contents

### From

The minimum value in the range.

Type: Integer

Required: No

### Step

The step value for the range. For example, if you have a range of 5,000 to 10,000, with a step value of 1,000, the valid values start at 5,000 and step up by 1,000. Even though 7,500 is within the range, it isn't a valid value for the range. The valid values are 5,000, 6,000, 7,000, 8,000...

Type: Integer

Required: No

### To

The maximum value in the range.

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](#)

# ResourcePendingMaintenanceActions

Describes the pending maintenance actions for a resource.

## Contents

### PendingMaintenanceActionDetails.PendingMaintenanceAction.N

A list that provides details about the pending maintenance actions for the resource.

Type: Array of [PendingMaintenanceAction](#) objects

Required: No

### ResourceIdentifier

The ARN of the resource that has pending maintenance actions.

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](#)

# ServerlessV2ScalingConfiguration

Contains the scaling configuration of a Neptune Serverless DB cluster.

For more information, see [Using Amazon Neptune Serverless](#) in the *Amazon Neptune User Guide*.

## Contents

### MaxCapacity

The maximum number of Neptune capacity units (NCUs) for a DB instance in a Neptune Serverless cluster. You can specify NCU values in half-step increments, such as 40, 40.5, 41, and so on.

Type: Double

Required: No

### MinCapacity

The minimum number of Neptune capacity units (NCUs) for a DB instance in a Neptune Serverless cluster. You can specify NCU values in half-step increments, such as 8, 8.5, 9, and so on.

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](#)

# ServerlessV2ScalingConfigurationInfo

Shows the scaling configuration for a Neptune Serverless DB cluster.

For more information, see [Using Amazon Neptune Serverless](#) in the *Amazon Neptune User Guide*.

## Contents

### MaxCapacity

The maximum number of Neptune capacity units (NCUs) for a DB instance in a Neptune Serverless cluster. You can specify NCU values in half-step increments, such as 40, 40.5, 41, and so on.

Type: Double

Required: No

### MinCapacity

The minimum number of Neptune capacity units (NCUs) for a DB instance in a Neptune Serverless cluster. You can specify NCU values in half-step increments, such as 8, 8.5, 9, and so on.

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](#)

# Subnet

Specifies a subnet.

This data type is used as a response element in the [DescribeDBSubnetGroups](#) action.

## Contents

### SubnetAvailabilityZone

Specifies the EC2 Availability Zone that the subnet is in.

Type: [AvailabilityZone](#) object

Required: No

### SubnetIdentifier

Specifies the identifier of the subnet.

Type: String

Required: No

### SubnetStatus

Specifies the status 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

Metadata assigned to an Amazon Neptune resource consisting of a key-value pair.

## Contents

### Key

A key is the required name of the tag. The string value can be from 1 to 128 Unicode characters in length and can't be prefixed with `aws:` or `ids:`. The string can only contain the set of Unicode letters, digits, white-space, `'_'`, `'!'`, `'/'`, `'='`, `'+'`, `'-'` (Java regex: `"^([\p{L}\p{Z}\p{N}_:/=+\-]*)$"`).

Type: String

Required: No

### Value

A value is the optional value of the tag. The string value can be from 1 to 256 Unicode characters in length and can't be prefixed with `aws:` or `ids:`. The string can only contain the set of Unicode letters, digits, white-space, `'_'`, `'!'`, `'/'`, `'='`, `'+'`, `'-'` (Java regex: `"^([\p{L}\p{Z}\p{N}_:/=+\-]*)$"`).

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](#)

# Timezone

A time zone associated with a [DBInstance](#).

## Contents

### TimezoneName

The name of the time 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](#)

# UpgradeTarget

The version of the database engine that a DB instance can be upgraded to.

## Contents

### AutoUpgrade

A value that indicates whether the target version is applied to any source DB instances that have `AutoMinorVersionUpgrade` set to true.

Type: Boolean

Required: No

### Description

The version of the database engine that a DB instance can be upgraded to.

Type: String

Required: No

### Engine

The name of the upgrade target database engine.

Type: String

Required: No

### EngineVersion

The version number of the upgrade target database engine.

Type: String

Required: No

### IsMajorVersionUpgrade

A value that indicates whether a database engine is upgraded to a major version.

Type: Boolean

Required: No



## SupportsGlobalDatabases

A value that indicates whether you can use Neptune global databases with the target engine version.

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](#)

# ValidDBInstanceModificationsMessage

Information about valid modifications that you can make to your DB instance. Contains the result of a successful call to the [DescribeValidDBInstanceModifications](#) action. You can use this information when you call [ModifyDBInstance](#).

## Contents

### Storage.ValidStorageOptions.N

Valid storage options for your DB instance.

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

# ValidStorageOptions

Not applicable. In Neptune the storage type is managed at the DB Cluster level.

## Contents

### IopsToStorageRatio.DoubleRange.N

Not applicable. In Neptune the storage type is managed at the DB Cluster level.

Type: Array of [DoubleRange](#) objects

Required: No

### ProvisionedIops.Range.N

Not applicable. In Neptune the storage type is managed at the DB Cluster level.

Type: Array of [Range](#) objects

Required: No

### StorageSize.Range.N

Not applicable. In Neptune the storage type is managed at the DB Cluster level.

Type: Array of [Range](#) objects

Required: No

### StorageType

Not applicable. In Neptune the storage type is managed at the DB Cluster level.

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](#)

# VpcSecurityGroupMembership

This data type is used as a response element for queries on VPC security group membership.

## Contents

### Status

The status of the VPC security group.

Type: String

Required: No

### VpcSecurityGroupId

The name of the VPC 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](#)

# 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