
Amazon FSx

API Reference



Amazon FSx: API Reference

Copyright © 2020 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

Amazon FSx API Reference	1
API Endpoint	1
API Version	1
Related Topics	2
Amazon FSx for Windows File Server Forums	2
Actions	3
CancelDataRepositoryTask	4
Request Syntax	4
Request Parameters	4
Response Syntax	4
Response Elements	4
Errors	5
Examples	5
See Also	6
CreateBackup	7
Request Syntax	7
Request Parameters	7
Response Syntax	8
Response Elements	10
Errors	10
See Also	11
CreateDataRepositoryTask	12
Request Syntax	12
Request Parameters	12
Response Syntax	13
Response Elements	14
Errors	14
Examples	15
See Also	16
CreateFileSystem	17
Request Syntax	17
Request Parameters	18
Response Syntax	20
Response Elements	21
Errors	22
See Also	23
CreateFileSystemFromBackup	24
Request Syntax	24
Request Parameters	25
Response Syntax	27
Response Elements	28
Errors	28
See Also	29
DeleteBackup	30
Request Syntax	30
Request Parameters	30
Response Syntax	30
Response Elements	31
Errors	31
See Also	32
DeleteFileSystem	33
Request Syntax	33
Request Parameters	33
Response Syntax	34

Response Elements	35
Errors	35
See Also	36
DescribeBackups	37
Request Syntax	37
Request Parameters	37
Response Syntax	38
Response Elements	40
Errors	40
See Also	41
DescribeDataRepositoryTasks	42
Request Syntax	42
Request Parameters	42
Response Syntax	43
Response Elements	44
Errors	44
Examples	44
See Also	45
DescribeFileSystems	47
Request Syntax	47
Request Parameters	47
Response Syntax	48
Response Elements	49
Errors	50
See Also	50
ListTagsForResource	51
Request Syntax	51
Request Parameters	51
Response Syntax	52
Response Elements	52
Errors	52
See Also	53
TagResource	54
Request Syntax	54
Request Parameters	54
Response Elements	54
Errors	54
See Also	55
UntagResource	56
Request Syntax	56
Request Parameters	56
Response Elements	56
Errors	56
See Also	57
UpdateFileSystem	58
Request Syntax	58
Request Parameters	58
Response Syntax	59
Response Elements	61
Errors	61
See Also	62
Data Types	63
ActiveDirectoryBackupAttributes	64
Contents	64
See Also	64
AdministrativeAction	65
Contents	65

See Also	66
AdministrativeActionFailureDetails	67
Contents	67
See Also	67
Backup	68
Contents	68
See Also	69
BackupFailureDetails	71
Contents	71
See Also	71
CompletionReport	72
Contents	72
See Also	73
CreateFileSystemLustreConfiguration	74
Contents	74
See Also	77
CreateFileSystemWindowsConfiguration	78
Contents	78
See Also	80
DataRepositoryConfiguration	81
Contents	81
See Also	82
DataRepositoryFailureDetails	83
Contents	83
See Also	83
DataRepositoryTask	84
Contents	84
See Also	86
DataRepositoryTaskFailureDetails	87
Contents	87
See Also	87
DataRepositoryTaskFilter	88
Contents	88
See Also	88
DataRepositoryTaskStatus	89
Contents	89
See Also	89
DeleteFileSystemLustreConfiguration	90
Contents	90
See Also	90
DeleteFileSystemLustreResponse	91
Contents	91
See Also	91
DeleteFileSystemWindowsConfiguration	92
Contents	92
See Also	92
DeleteFileSystemWindowsResponse	93
Contents	93
See Also	93
FileSystem	94
Contents	94
See Also	97
FileSystemFailureDetails	98
Contents	98
See Also	98
Filter	99
Contents	99

See Also	99
LustreFileSystemConfiguration	100
Contents	100
See Also	102
SelfManagedActiveDirectoryAttributes	103
Contents	103
See Also	104
SelfManagedActiveDirectoryConfiguration	105
Contents	105
See Also	106
SelfManagedActiveDirectoryConfigurationUpdates	107
Contents	107
See Also	107
Tag	109
Contents	109
See Also	109
UpdateFileSystemLustreConfiguration	110
Contents	110
See Also	111
UpdateFileSystemWindowsConfiguration	112
Contents	112
See Also	113
WindowsFileSystemConfiguration	114
Contents	114
See Also	116
Common Parameters	117
Common Errors	119

Amazon FSx API Reference

The Amazon FSx API is a network protocol based on [HTTP \(RFC 2616\)](#). For each API call, you make an HTTP request to the region-specific Amazon FSx API endpoint for the AWS Region where you want to manage file systems. The API uses JSON (RFC 4627) documents for HTTP request/response bodies.

The Amazon FSx API is an RPC model. In this model, there is a fixed set of operations and the syntax for each operation is known to clients without any prior interaction. In the following section, you can find a description of each API operation using an abstract RPC notation. Each has an operation name that doesn't appear on the wire. For each operation, the topic specifies the mapping to HTTP request elements.

The specific Amazon FSx operation to which a given request maps is determined by a combination of the request's method (GET, PUT, POST, or DELETE) and which of the various patterns its Request-URI matches. If the operation is PUT or POST, Amazon FSx extracts call arguments from the Request-URI path segment, query parameters, and the JSON object in the request body.

Although operation names, such as `CreateFileSystem`, don't appear on the wire, these names are meaningful in AWS Identity and Access Management (IAM) policies. The operation name is also used to name commands in command-line tools and elements of the AWS SDKs. For example, there is a AWS CLI command named `create-file-system` that maps to the `CreateFileSystem` operation. The operation name also appears in AWS CloudTrail logs for Amazon FSx API calls.

API Endpoint

The API endpoint is the DNS name used as a host in the HTTP URI for the API calls. These API endpoints are specific to AWS Regions and take the following form.

```
fsx.aws-region.amazonaws.com
```

For example, the Amazon FSx API endpoint for the US East (N. Virginia) Region is the following.

```
fsx.us-east-1.amazonaws.com
```

For a list of AWS Regions that Amazon FSx supports (where you can create and manage file systems), see [Amazon FSx](#) in the *AWS General Reference*.

The region-specific API endpoint defines the scope of the Amazon FSx resources that are accessible when you make an API call. For example, when you call the `DescribeFileSystems` operation using the preceding endpoint, you get a list of file systems in the US West (Oregon) Region that have been created in your account.

API Version

The version of the API being used for a call is identified by the first path segment of the request URI, and its form is an ISO 8601 date. The documentation describes API version 2018-03-01.

Related Topics

The following related content provide information on the necessary permissions for these API operations using IAM policies.

- [Amazon FSx for Windows File Server API Permissions: Actions, Resources, and Conditions Reference](#) in the Amazon FSx for Windows File Sever User Guide
- [Amazon FSx for Lustre API Permissions: Actions, Resources, and Conditions Reference](#) in the Amazon FSx for Lustre User Guide

Amazon FSx for Windows File Server Forums

If you encounter issues while using Amazon FSx use the forums:

- [Amazon FSx for Windows File Server forums.](#)
- [Amazon FSx for Lustre forums.](#)

Actions

The following actions are supported:

- [CancelDataRepositoryTask](#) (p. 4)
- [CreateBackup](#) (p. 7)
- [CreateDataRepositoryTask](#) (p. 12)
- [CreateFileSystem](#) (p. 17)
- [CreateFileSystemFromBackup](#) (p. 24)
- [DeleteBackup](#) (p. 30)
- [DeleteFileSystem](#) (p. 33)
- [DescribeBackups](#) (p. 37)
- [DescribeDataRepositoryTasks](#) (p. 42)
- [DescribeFileSystems](#) (p. 47)
- [ListTagsForResource](#) (p. 51)
- [TagResource](#) (p. 54)
- [UntagResource](#) (p. 56)
- [UpdateFileSystem](#) (p. 58)

CancelDataRepositoryTask

Cancels an existing Amazon FSx for Lustre data repository task if that task is in either the `PENDING` or `EXECUTING` state. When you cancel a task, Amazon FSx does the following.

- Any files that FSx has already exported are not reverted.
- FSx continues to export any files that are "in-flight" when the cancel operation is received.
- FSx does not export any files that have not yet been exported.

Request Syntax

```
{  
  "TaskId": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 117\)](#).

The request accepts the following data in JSON format.

TaskId (p. 4)

Specifies the data repository task to cancel.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 128.

Pattern: `^(task-[0-9a-f]{17,})$`

Required: Yes

Response Syntax

```
{  
  "Lifecycle": "string",  
  "TaskId": "string"  
}
```

Response Elements

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

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

Lifecycle (p. 4)

The lifecycle status of the data repository task, as follows:

- `PENDING` - Amazon FSx has not started the task.
- `EXECUTING` - Amazon FSx is processing the task.

- **FAILED** - Amazon FSx was not able to complete the task. For example, there may be files the task failed to process. The [DataRepositoryTaskFailureDetails \(p. 87\)](#) property provides more information about task failures.
- **SUCCEEDED** - FSx completed the task successfully.
- **CANCELED** - Amazon FSx canceled the task and it did not complete.
- **CANCELING** - FSx is in process of canceling the task.

Type: String

Valid Values: **PENDING** | **EXECUTING** | **FAILED** | **SUCCEEDED** | **CANCELED** | **CANCELING**

TaskId (p. 4)

The ID of the task being canceled.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 128.

Pattern: `^(task-[0-9a-f]{17,})$`

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 119\)](#).

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

DataRepositoryTaskEnded

The data repository task could not be canceled because the task has already ended.

HTTP Status Code: 400

DataRepositoryTaskNotFound

The data repository task or tasks you specified could not be found.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

UnsupportedOperation

The requested operation is not supported for this resource or API.

HTTP Status Code: 400

Examples

Cancel a Data Repository Task

The following request cancels a specific data repository task by using the `TaskId` request parameter.

Sample Request

```
POST /2015-02-01/cancel-data-repository-task HTTP/1.1

{
  "TaskId": ["task-0123456789abcdef0"]
}
```

Sample Response

```
HTTP/1.1 200 success
x-amzn-RequestId: 12345678-1234-abcd-5678-0123456789abc

{
  "Status": "CANCELING",
  "TaskId": "task-0123456789abcdef0"
}
```

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

CreateBackup

Creates a backup of an existing Amazon FSx file system. Creating regular backups for your file system is a best practice, enabling you to restore a file system from a backup if an issue arises with the original file system.

For Amazon FSx for Lustre file systems, you can create a backup only for file systems with the following configuration:

- a Persistent deployment type
- is *not* linked to a data repository.

For more information about backing up Amazon FSx for Lustre file systems, see [Working with FSx for Lustre backups](#).

For more information about backing up Amazon FSx for Lustre file systems, see [Working with FSx for Windows backups](#).

If a backup with the specified client request token exists, and the parameters match, this operation returns the description of the existing backup. If a backup specified client request token exists, and the parameters don't match, this operation returns `IncompatibleParameterError`. If a backup with the specified client request token doesn't exist, `CreateBackup` does the following:

- Creates a new Amazon FSx backup with an assigned ID, and an initial lifecycle state of `CREATING`.
- Returns the description of the backup.

By using the idempotent operation, you can retry a `CreateBackup` operation without the risk of creating an extra backup. This approach can be useful when an initial call fails in a way that makes it unclear whether a backup was created. If you use the same client request token and the initial call created a backup, the operation returns a successful result because all the parameters are the same.

The `CreateBackup` operation returns while the backup's lifecycle state is still `CREATING`. You can check the backup creation status by calling the [DescribeBackups \(p. 37\)](#) operation, which returns the backup state along with other information.

Request Syntax

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

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 117\)](#).

The request accepts the following data in JSON format.

ClientRequestToken (p. 7)

(Optional) A string of up to 64 ASCII characters that Amazon FSx uses to ensure idempotent creation. This string is automatically filled on your behalf when you use the AWS Command Line Interface (AWS CLI) or an AWS SDK.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[A-Za-z0-9_.-]{0,63}`\$

Required: No

FileSystemId (p. 7)

The ID of the file system to back up.

Type: String

Length Constraints: Minimum length of 11. Maximum length of 21.

Pattern: `^(fs-[0-9a-f]{8,})$`

Required: Yes

Tags (p. 7)

(Optional) The tags to apply to the backup at backup creation. The key value of the Name tag appears in the console as the backup name. If you have set `CopyTagsToBackups` to true, and you specify one or more tags using the `CreateBackup` action, no existing file system tags are copied from the file system to the backup.

Type: Array of [Tag \(p. 109\)](#) objects

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

Required: No

Response Syntax

```
{
  "Backup": {
    "BackupId": "string",
    "CreationTime": number,
    "DirectoryInformation": {
      "ActiveDirectoryId": "string",
      "DomainName": "string"
    },
    "FailureDetails": {
      "Message": "string"
    },
    "FileSystem": {
      "AdministrativeActions": [
        {
          "AdministrativeActionType": "string",
          "FailureDetails": {
            "Message": "string"
          },
          "ProgressPercent": number,
          "RequestTime": number,
          "Status": "string",

```

```

        "TargetFileSystemValues": "FileSystem"
    }
],
"CreationTime": number,
"DNSName": "string",
"FailureDetails": {
    "Message": "string"
},
"FileSystemId": "string",
"FileSystemType": "string",
"KmsKeyId": "string",
"Lifecycle": "string",
"LustreConfiguration": {
    "AutomaticBackupRetentionDays": number,
    "CopyTagsToBackups": boolean,
    "DailyAutomaticBackupStartTime": "string",
    "DataRepositoryConfiguration": {
        "AutoImportPolicy": "string",
        "ExportPath": "string",
        "FailureDetails": {
            "Message": "string"
        },
        "ImportedFileChunkSize": number,
        "ImportPath": "string",
        "Lifecycle": "string"
    },
    "DeploymentType": "string",
    "DriveCacheType": "string",
    "MountName": "string",
    "PerUnitStorageThroughput": number,
    "WeeklyMaintenanceStartTime": "string"
},
"NetworkInterfaceIds": [ "string" ],
"OwnerId": "string",
"ResourceARN": "string",
"StorageCapacity": number,
"StorageType": "string",
"SubnetIds": [ "string" ],
"Tags": [
    {
        "Key": "string",
        "Value": "string"
    }
],
"VpcId": "string",
"WindowsConfiguration": {
    "ActiveDirectoryId": "string",
    "AutomaticBackupRetentionDays": number,
    "CopyTagsToBackups": boolean,
    "DailyAutomaticBackupStartTime": "string",
    "DeploymentType": "string",
    "MaintenanceOperationsInProgress": [ "string" ],
    "PreferredFileServerIp": "string",
    "PreferredSubnetId": "string",
    "RemoteAdministrationEndpoint": "string",
    "SelfManagedActiveDirectoryConfiguration": {
        "DnsIps": [ "string" ],
        "DomainName": "string",
        "FileSystemAdministratorsGroup": "string",
        "OrganizationalUnitDistinguishedName": "string",
        "UserName": "string"
    },
    "ThroughputCapacity": number,
    "WeeklyMaintenanceStartTime": "string"
}
},
}

```

```
    "KmsKeyId": "string",
    "Lifecycle": "string",
    "ProgressPercent": number,
    "ResourceARN": "string",
    "Tags": [
      {
        "Key": "string",
        "Value": "string"
      }
    ],
    "Type": "string"
  }
}
```

Response Elements

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

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

Backup (p. 8)

A description of the backup.

Type: [Backup \(p. 68\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 119\)](#).

BackupInProgress

Another backup is already under way. Wait for completion before initiating additional backups of this file system.

HTTP Status Code: 400

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

FileSystemNotFound

No Amazon FSx file systems were found based upon supplied parameters.

HTTP Status Code: 400

IncompatibleParameterError

The error returned when a second request is received with the same client request token but different parameters settings. A client request token should always uniquely identify a single request.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

ServiceLimitExceeded

An error indicating that a particular service limit was exceeded. You can increase some service limits by contacting AWS Support.

HTTP Status Code: 400

UnsupportedOperation

The requested operation is not supported for this resource or API.

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

CreateDataRepositoryTask

Creates an Amazon FSx for Lustre data repository task. You use data repository tasks to perform bulk operations between your Amazon FSx file system and its linked data repository. An example of a data repository task is exporting any data and metadata changes, including POSIX metadata, to files, directories, and symbolic links (symlinks) from your FSx file system to its linked data repository. A `CreateDataRepositoryTask` operation will fail if a data repository is not linked to the FSx file system. To learn more about data repository tasks, see [Using Data Repository Tasks](#). To learn more about linking a data repository to your file system, see [Setting the Export Prefix](#).

Request Syntax

```
{
  "ClientRequestToken": "string",
  "FileSystemId": "string",
  "Paths": [ "string" ],
  "Report": {
    "Enabled": boolean,
    "Format": "string",
    "Path": "string",
    "Scope": "string"
  },
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ],
  "Type": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 117).

The request accepts the following data in JSON format.

[ClientRequestToken](#) (p. 12)

(Optional) An idempotency token for resource creation, in a string of up to 64 ASCII characters. This token is automatically filled on your behalf when you use the AWS Command Line Interface (AWS CLI) or an AWS SDK.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[A-Za-z0-9_.-]{0,63}`\$

Required: No

[FileSystemId](#) (p. 12)

The globally unique ID of the file system, assigned by Amazon FSx.

Type: String

Length Constraints: Minimum length of 11. Maximum length of 21.

Pattern: `^(fs-[0-9a-f]{8,})$`

Required: Yes

Paths (p. 12)

(Optional) The path or paths on the Amazon FSx file system to use when the data repository task is processed. The default path is the file system root directory. The paths you provide need to be relative to the mount point of the file system. If the mount point is `/mnt/fsx` and `/mnt/fsx/path1` is a directory or file on the file system you want to export, then the path to provide is `path1`. If a path that you provide isn't valid, the task fails.

Type: Array of strings

Array Members: Maximum number of 100 items.

Length Constraints: Minimum length of 0. Maximum length of 4096.

Pattern: `^[^\u0000\u0085\u2028\u2029\x\n]{0,4096}$`

Required: No

Report (p. 12)

Defines whether or not Amazon FSx provides a `CompletionReport` once the task has completed. A `CompletionReport` provides a detailed report on the files that Amazon FSx processed that meet the criteria specified by the `Scope` parameter. For more information, see [Working with Task Completion Reports](#).

Type: [CompletionReport \(p. 72\)](#) object

Required: Yes

Tags (p. 12)

A list of `Tag` values, with a maximum of 50 elements.

Type: Array of [Tag \(p. 109\)](#) objects

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

Required: No

Type (p. 12)

Specifies the type of data repository task to create.

Type: String

Valid Values: `EXPORT_TO_REPOSITORY`

Required: Yes

Response Syntax

```
{
  "DataRepositoryTask": {
    "CreationTime": number,
    "EndTime": number,
    "FailureDetails": {
      "Message": "string"
    },
    "FileSystemId": "string",
```

```
"Lifecycle": "string",
"Paths": [ "string" ],
"Report": {
  "Enabled": boolean,
  "Format": "string",
  "Path": "string",
  "Scope": "string"
},
"ResourceARN": "string",
"StartTime": number,
"Status": {
  "FailedCount": number,
  "LastUpdatedTime": number,
  "SucceededCount": number,
  "TotalCount": number
},
"Tags": [
  {
    "Key": "string",
    "Value": "string"
  }
],
"TaskId": "string",
"Type": "string"
}
```

Response Elements

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

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

DataRepositoryTask (p. 13)

The description of the data repository task that you just created.

Type: [DataRepositoryTask \(p. 84\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 119\)](#).

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

DataRepositoryTaskExecuting

An existing data repository task is currently executing on the file system. Wait until the existing task has completed, then create the new task.

HTTP Status Code: 400

FileSystemNotFound

No Amazon FSx file systems were found based upon supplied parameters.

HTTP Status Code: 400

IncompatibleParameterError

The error returned when a second request is received with the same client request token but different parameters settings. A client request token should always uniquely identify a single request.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

ServiceLimitExceeded

An error indicating that a particular service limit was exceeded. You can increase some service limits by contacting AWS Support.

HTTP Status Code: 400

UnsupportedOperation

The requested operation is not supported for this resource or API.

HTTP Status Code: 400

Examples

Create a Data Repository Task

The following request creates an `EXPORT_TO_REPOSITORY` data repository task for the specified file system.

Sample Request

```
POST /2015-02-01/create-data-repository-task HTTP/1.1
Host: fsx.us-east-1.amazonaws.com
x-amz-date: 20140620T221118Z
Authorization: <...>
Content-Type: application/json
Content-Length: 160

{
  "FileSystemId": "fs-0123456789abcdef0",
  "Type": "EXPORT_TO_REPOSITORY",
  "Paths": ["path1", "path2/file1"],
  "Report": {
    "Enabled": true,
    "Path": "s3://myBucket/FSxLustre20191118T225838Z/myreports",
    "Format": "REPORT_CSV_20191124",
    "Scope": "FAILED_FILES_ONLY"
  },
}
```

Sample Response

```
HTTP/1.1 200 success
x-amzn-RequestId: c3616af3-33fa-40ad-ae0d-d3895a2c3a1f
```

```
{
  "Task": {
    "TaskId": "task-0123456789abcdef1",
    "TaskType": "EXPORT_TO_REPOSITORY",
    "Lifecycle": "PENDING",
    "FileSystemId": "fs-0123456789abcdef0",
    "Paths": ["path1", "path2/file1"],
    "TaskReport": {
      "Path": "s3://myBucket/FSxLustre20191118T225838Z/myreports",
      "Format": "REPORT_CSV_20191124",
      "Enabled": true,
      "Scope": "FAILED_FILES_ONLY"
    },
    "Tags": [],
    "CreationTime": "2018-12-17T18:18:18.000Z",
    "ClientRequestToken": "10192019-drt-12",
    "ResourceARN": "arn:aws:fsx:us-east-1:123456789012:task:task-123f8cd8e330c1321"
  }
}
```

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

CreateFileSystem

Creates a new, empty Amazon FSx file system.

If a file system with the specified client request token exists and the parameters match, `CreateFileSystem` returns the description of the existing file system. If a file system specified client request token exists and the parameters don't match, this call returns `IncompatibleParameterError`. If a file system with the specified client request token doesn't exist, `CreateFileSystem` does the following:

- Creates a new, empty Amazon FSx file system with an assigned ID, and an initial lifecycle state of `CREATING`.
- Returns the description of the file system.

This operation requires a client request token in the request that Amazon FSx uses to ensure idempotent creation. This means that calling the operation multiple times with the same client request token has no effect. By using the idempotent operation, you can retry a `CreateFileSystem` operation without the risk of creating an extra file system. This approach can be useful when an initial call fails in a way that makes it unclear whether a file system was created. Examples are if a transport level timeout occurred, or your connection was reset. If you use the same client request token and the initial call created a file system, the client receives success as long as the parameters are the same.

Note

The `CreateFileSystem` call returns while the file system's lifecycle state is still `CREATING`. You can check the file-system creation status by calling the [DescribeFileSystems](#) (p. 47) operation, which returns the file system state along with other information.

Request Syntax

```
{
  "ClientRequestToken": "string",
  "FileSystemType": "string",
  "KmsKeyId": "string",
  "LustreConfiguration": {
    "AutoImportPolicy": "string",
    "AutomaticBackupRetentionDays": number,
    "CopyTagsToBackups": boolean,
    "DailyAutomaticBackupStartTime": "string",
    "DeploymentType": "string",
    "DriveCacheType": "string",
    "ExportPath": "string",
    "ImportedFileChunkSize": number,
    "ImportPath": "string",
    "PerUnitStorageThroughput": number,
    "WeeklyMaintenanceStartTime": "string"
  },
  "SecurityGroupIds": [ "string" ],
  "StorageCapacity": number,
  "StorageType": "string",
  "SubnetIds": [ "string" ],
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ],
  "WindowsConfiguration": {
    "ActiveDirectoryId": "string",
    "AutomaticBackupRetentionDays": number,
```

```
"CopyTagsToBackups": boolean,  
"DailyAutomaticBackupStartTime": "string",  
"DeploymentType": "string",  
"PreferredSubnetId": "string",  
"SelfManagedActiveDirectoryConfiguration": {  
  "DnsIps": [ "string" ],  
  "DomainName": "string",  
  "FileSystemAdministratorsGroup": "string",  
  "OrganizationalUnitDistinguishedName": "string",  
  "Password": "string",  
  "UserName": "string"  
},  
"ThroughputCapacity": number,  
"WeeklyMaintenanceStartTime": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 117\)](#).

The request accepts the following data in JSON format.

[ClientRequestToken \(p. 17\)](#)

A string of up to 64 ASCII characters that Amazon FSx uses to ensure idempotent creation. This string is automatically filled on your behalf when you use the AWS Command Line Interface (AWS CLI) or an AWS SDK.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[A-Za-z0-9_.-]{0,63}$`

Required: No

[FileSystemType \(p. 17\)](#)

The type of Amazon FSx file system to create, either `WINDOWS` or `LUSTRE`.

Type: String

Valid Values: `WINDOWS` | `LUSTRE`

Required: Yes

[KmsKeyId \(p. 17\)](#)

The ID of the AWS Key Management Service (AWS KMS) key used to encrypt the file system's data for Amazon FSx for Windows File Server file systems and Amazon FSx for Lustre `PERSISTENT_1` file systems at rest. In either case, if not specified, the Amazon FSx managed key is used. The Amazon FSx for Lustre `SCRATCH_1` and `SCRATCH_2` file systems are always encrypted at rest using Amazon FSx managed keys. For more information, see [Encrypt](#) in the *AWS Key Management Service API Reference*.

Type: String

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

Pattern: `^\.{1,2048}$`

Required: No

LustreConfiguration (p. 17)

The Lustre configuration for the file system being created.

Type: [CreateFileSystemLustreConfiguration \(p. 74\)](#) object

Required: No

SecurityGroupIds (p. 17)

A list of IDs specifying the security groups to apply to all network interfaces created for file system access. This list isn't returned in later requests to describe the file system.

Type: Array of strings

Array Members: Maximum number of 50 items.

Length Constraints: Minimum length of 11. Maximum length of 20.

Pattern: `^(sg-[0-9a-f]{8,})$`

Required: No

StorageCapacity (p. 17)

Sets the storage capacity of the file system that you're creating.

For Lustre file systems:

- For `SCRATCH_2` and `PERSISTENT_1` SSD deployment types, valid values are 1200 GiB, 2400 GiB, and increments of 2400 GiB.
- For `PERSISTENT` HDD file systems, valid values are increments of 6000 GiB for 12 MB/s/TiB file systems and increments of 1800 GiB for 40 MB/s/TiB file systems.
- For `SCRATCH_1` deployment type, valid values are 1200 GiB, 2400 GiB, and increments of 3600 GiB.

For Windows file systems:

- If `StorageType=SSD`, valid values are 32 GiB - 65,536 GiB (64 TiB).
- If `StorageType=HDD`, valid values are 2000 GiB - 65,536 GiB (64 TiB).

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 2147483647.

Required: Yes

StorageType (p. 17)

Sets the storage type for the file system you're creating. Valid values are `SSD` and `HDD`.

- Set to `SSD` to use solid state drive storage. SSD is supported on all Windows and Lustre deployment types.
- Set to `HDD` to use hard disk drive storage. HDD is supported on `SINGLE_AZ_2` and `MULTI_AZ_1` Windows file system deployment types, and on `PERSISTENT` Lustre file system deployment types.

Default value is `SSD`. For more information, see [Storage Type Options](#) in the *Amazon FSx for Windows User Guide* and [Multiple Storage Options](#) in the *Amazon FSx for Lustre User Guide*.

Type: String

Valid Values: `SSD` | `HDD`

Required: No

SubnetIds (p. 17)

Specifies the IDs of the subnets that the file system will be accessible from. For Windows `MULTI_AZ_1` file system deployment types, provide exactly two subnet IDs, one for the preferred file server and one for the standby file server. You specify one of these subnets as the preferred subnet using the `WindowsConfiguration > PreferredSubnetID` property.

For Windows `SINGLE_AZ_1` and `SINGLE_AZ_2` file system deployment types and Lustre file systems, provide exactly one subnet ID. The file server is launched in that subnet's Availability Zone.

Type: Array of strings

Array Members: Maximum number of 50 items.

Length Constraints: Minimum length of 15. Maximum length of 24.

Pattern: `^(subnet-[0-9a-f]{8,})$`

Required: Yes

Tags (p. 17)

The tags to apply to the file system being created. The key value of the `Name` tag appears in the console as the file system name.

Type: Array of [Tag \(p. 109\)](#) objects

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

Required: No

WindowsConfiguration (p. 17)

The Microsoft Windows configuration for the file system being created.

Type: [CreateFileSystemWindowsConfiguration \(p. 78\)](#) object

Required: No

Response Syntax

```
{
  "FileSystem": {
    "AdministrativeActions": [
      {
        "AdministrativeActionType": "string",
        "FailureDetails": {
          "Message": "string"
        },
        "ProgressPercent": number,
        "RequestTime": number,
        "Status": "string",
        "TargetFileSystemValues": "FileSystem"
      }
    ],
    "CreationTime": number,
    "DNSName": "string",
    "FailureDetails": {
      "Message": "string"
    },
    "FileSystemId": "string",
```

```

"FileSystemType": "string",
"KmsKeyId": "string",
"Lifecycle": "string",
"LustreConfiguration": {
  "AutomaticBackupRetentionDays": number,
  "CopyTagsToBackups": boolean,
  "DailyAutomaticBackupStartTime": "string",
  "DataRepositoryConfiguration": {
    "AutoImportPolicy": "string",
    "ExportPath": "string",
    "FailureDetails": {
      "Message": "string"
    },
    "ImportedFileChunkSize": number,
    "ImportPath": "string",
    "Lifecycle": "string"
  },
  "DeploymentType": "string",
  "DriveCacheType": "string",
  "MountName": "string",
  "PerUnitStorageThroughput": number,
  "WeeklyMaintenanceStartTime": "string"
},
"NetworkInterfaceIds": [ "string" ],
"OwnerId": "string",
"ResourceARN": "string",
"StorageCapacity": number,
"StorageType": "string",
"SubnetIds": [ "string" ],
"Tags": [
  {
    "Key": "string",
    "Value": "string"
  }
],
"VpcId": "string",
"WindowsConfiguration": {
  "ActiveDirectoryId": "string",
  "AutomaticBackupRetentionDays": number,
  "CopyTagsToBackups": boolean,
  "DailyAutomaticBackupStartTime": "string",
  "DeploymentType": "string",
  "MaintenanceOperationsInProgress": [ "string" ],
  "PreferredFileServerIp": "string",
  "PreferredSubnetId": "string",
  "RemoteAdministrationEndpoint": "string",
  "SelfManagedActiveDirectoryConfiguration": {
    "DnsIps": [ "string" ],
    "DomainName": "string",
    "FileSystemAdministratorsGroup": "string",
    "OrganizationalUnitDistinguishedName": "string",
    "UserName": "string"
  },
  "ThroughputCapacity": number,
  "WeeklyMaintenanceStartTime": "string"
}
}
}

```

Response Elements

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

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

[FileSystem \(p. 20\)](#)

The configuration of the file system that was created.

Type: [FileSystem \(p. 94\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 119\)](#).

ActiveDirectoryError

An Active Directory error.

HTTP Status Code: 400

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

IncompatibleParameterError

The error returned when a second request is received with the same client request token but different parameters settings. A client request token should always uniquely identify a single request.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

InvalidExportPath

The path provided for data repository export isn't valid.

HTTP Status Code: 400

InvalidImportPath

The path provided for data repository import isn't valid.

HTTP Status Code: 400

InvalidNetworkSettings

One or more network settings specified in the request are invalid. `InvalidVpcId` means that the ID passed for the virtual private cloud (VPC) is invalid. `InvalidSubnetIds` returns the list of IDs for subnets that are either invalid or not part of the VPC specified. `InvalidSecurityGroupIds` returns the list of IDs for security groups that are either invalid or not part of the VPC specified.

HTTP Status Code: 400

InvalidPerUnitStorageThroughput

An invalid value for `PerUnitStorageThroughput` was provided. Please create your file system again, using a valid value.

HTTP Status Code: 400

MissingFileSystemConfiguration

A file system configuration is required for this operation.

HTTP Status Code: 400

ServiceLimitExceeded

An error indicating that a particular service limit was exceeded. You can increase some service limits by contacting AWS Support.

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

CreateFileSystemFromBackup

Creates a new Amazon FSx file system from an existing Amazon FSx backup.

If a file system with the specified client request token exists and the parameters match, this operation returns the description of the file system. If a client request token specified by the file system exists and the parameters don't match, this call returns `IncompatibleParameterError`. If a file system with the specified client request token doesn't exist, this operation does the following:

- Creates a new Amazon FSx file system from backup with an assigned ID, and an initial lifecycle state of `CREATING`.
- Returns the description of the file system.

Parameters like Active Directory, default share name, automatic backup, and backup settings default to the parameters of the file system that was backed up, unless overridden. You can explicitly supply other settings.

By using the idempotent operation, you can retry a `CreateFileSystemFromBackup` call without the risk of creating an extra file system. This approach can be useful when an initial call fails in a way that makes it unclear whether a file system was created. Examples are if a transport level timeout occurred, or your connection was reset. If you use the same client request token and the initial call created a file system, the client receives success as long as the parameters are the same.

Note

The `CreateFileSystemFromBackup` call returns while the file system's lifecycle state is still `CREATING`. You can check the file-system creation status by calling the [DescribeFileSystems \(p. 47\)](#) operation, which returns the file system state along with other information.

Request Syntax

```
{
  "BackupId": "string",
  "ClientRequestToken": "string",
  "LustreConfiguration": {
    "AutoImportPolicy": "string",
    "AutomaticBackupRetentionDays": number,
    "CopyTagsToBackups": boolean,
    "DailyAutomaticBackupStartTime": "string",
    "DeploymentType": "string",
    "DriveCacheType": "string",
    "ExportPath": "string",
    "ImportedFileChunkSize": number,
    "ImportPath": "string",
    "PerUnitStorageThroughput": number,
    "WeeklyMaintenanceStartTime": "string"
  },
  "SecurityGroupIds": [ "string" ],
  "StorageType": "string",
  "SubnetIds": [ "string" ],
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ],
  "WindowsConfiguration": {
    "ActiveDirectoryId": "string",
    "AutomaticBackupRetentionDays": number,
```

```
"CopyTagsToBackups": boolean,
"DailyAutomaticBackupStartTime": "string",
"DeploymentType": "string",
"PreferredSubnetId": "string",
"SelfManagedActiveDirectoryConfiguration": {
  "DnsIps": [ "string" ],
  "DomainName": "string",
  "FileSystemAdministratorsGroup": "string",
  "OrganizationalUnitDistinguishedName": "string",
  "Password": "string",
  "UserName": "string"
},
"ThroughputCapacity": number,
"WeeklyMaintenanceStartTime": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 117).

The request accepts the following data in JSON format.

[BackupId](#) (p. 24)

The ID of the backup. Specifies the backup to use if you're creating a file system from an existing backup.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 128.

Pattern: `^(backup-[0-9a-f]{8,})$`

Required: Yes

[ClientRequestToken](#) (p. 24)

A string of up to 64 ASCII characters that Amazon FSx uses to ensure idempotent creation. This string is automatically filled on your behalf when you use the AWS Command Line Interface (AWS CLI) or an AWS SDK.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[A-Za-z0-9_.-]{0,63}$`

Required: No

[LustreConfiguration](#) (p. 24)

The Lustre configuration for the file system being created.

Type: [CreateFileSystemLustreConfiguration](#) (p. 74) object

Required: No

[SecurityGroupIds](#) (p. 24)

A list of IDs for the security groups that apply to the specified network interfaces created for file system access. These security groups apply to all network interfaces. This value isn't returned in later [DescribeFileSystem](#) requests.

Type: Array of strings

Array Members: Maximum number of 50 items.

Length Constraints: Minimum length of 11. Maximum length of 20.

Pattern: `^(sg-[0-9a-f]{8,})$`

Required: No

StorageType (p. 24)

Sets the storage type for the Windows file system you're creating from a backup. Valid values are `SSD` and `HDD`.

- Set to `SSD` to use solid state drive storage. Supported on all Windows deployment types.
- Set to `HDD` to use hard disk drive storage. Supported on `SINGLE_AZ_2` and `MULTI_AZ_1` Windows file system deployment types.

Default value is `SSD`.

Note

`HDD` and `SSD` storage types have different minimum storage capacity requirements. A restored file system's storage capacity is tied to the file system that was backed up. You can create a file system that uses `HDD` storage from a backup of a file system that used `SSD` storage only if the original `SSD` file system had a storage capacity of at least 2000 GiB.

Type: String

Valid Values: `SSD` | `HDD`

Required: No

SubnetIds (p. 24)

Specifies the IDs of the subnets that the file system will be accessible from. For Windows `MULTI_AZ_1` file system deployment types, provide exactly two subnet IDs, one for the preferred file server and one for the standby file server. You specify one of these subnets as the preferred subnet using the `WindowsConfiguration > PreferredSubnetID` property.

For Windows `SINGLE_AZ_1` and `SINGLE_AZ_2` deployment types and Lustre file systems, provide exactly one subnet ID. The file server is launched in that subnet's Availability Zone.

Type: Array of strings

Array Members: Maximum number of 50 items.

Length Constraints: Minimum length of 15. Maximum length of 24.

Pattern: `^(subnet-[0-9a-f]{8,})$`

Required: Yes

Tags (p. 24)

The tags to be applied to the file system at file system creation. The key value of the `Name` tag appears in the console as the file system name.

Type: Array of [Tag \(p. 109\)](#) objects

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

Required: No

WindowsConfiguration (p. 24)

The configuration for this Microsoft Windows file system.

Type: [CreateFileSystemWindowsConfiguration \(p. 78\)](#) object

Required: No

Response Syntax

```
{
  "FileSystem": {
    "AdministrativeActions": [
      {
        "AdministrativeActionType": "string",
        "FailureDetails": {
          "Message": "string"
        },
        "ProgressPercent": number,
        "RequestTime": number,
        "Status": "string",
        "TargetFileSystemValues": "FileSystem"
      }
    ],
    "CreationTime": number,
    "DNSName": "string",
    "FailureDetails": {
      "Message": "string"
    },
    "FileSystemId": "string",
    "FileSystemType": "string",
    "KmsKeyId": "string",
    "Lifecycle": "string",
    "LustreConfiguration": {
      "AutomaticBackupRetentionDays": number,
      "CopyTagsToBackups": boolean,
      "DailyAutomaticBackupStartTime": "string",
      "DataRepositoryConfiguration": {
        "AutoImportPolicy": "string",
        "ExportPath": "string",
        "FailureDetails": {
          "Message": "string"
        },
        "ImportedFileChunkSize": number,
        "ImportPath": "string",
        "Lifecycle": "string"
      },
      "DeploymentType": "string",
      "DriveCacheType": "string",
      "MountName": "string",
      "PerUnitStorageThroughput": number,
      "WeeklyMaintenanceStartTime": "string"
    },
    "NetworkInterfaceIds": [ "string" ],
    "OwnerId": "string",
    "ResourceARN": "string",
    "StorageCapacity": number,
    "StorageType": "string",
    "SubnetIds": [ "string" ],
    "Tags": [
      {
        "Key": "string",
        "Value": "string"
      }
    ],
    "VpcId": "string",
    "WindowsConfiguration": {
```

```
"ActiveDirectoryId": "string",
"AutomaticBackupRetentionDays": number,
"CopyTagsToBackups": boolean,
"DailyAutomaticBackupStartTime": "string",
"DeploymentType": "string",
"MaintenanceOperationsInProgress": [ "string" ],
"PreferredFileServerIp": "string",
"PreferredSubnetId": "string",
"RemoteAdministrationEndpoint": "string",
"SelfManagedActiveDirectoryConfiguration": {
  "DnsIps": [ "string" ],
  "DomainName": "string",
  "FileSystemAdministratorsGroup": "string",
  "OrganizationalUnitDistinguishedName": "string",
  "UserName": "string"
},
"ThroughputCapacity": number,
"WeeklyMaintenanceStartTime": "string"
}
}
```

Response Elements

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

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

FileSystem (p. 27)

A description of the file system.

Type: [FileSystem \(p. 94\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 119\)](#).

ActiveDirectoryError

An Active Directory error.

HTTP Status Code: 400

BackupNotFound

No Amazon FSx backups were found based upon the supplied parameters.

HTTP Status Code: 400

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

IncompatibleParameterError

The error returned when a second request is received with the same client request token but different parameters settings. A client request token should always uniquely identify a single request.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

InvalidNetworkSettings

One or more network settings specified in the request are invalid. `InvalidVpcId` means that the ID passed for the virtual private cloud (VPC) is invalid. `InvalidSubnetIds` returns the list of IDs for subnets that are either invalid or not part of the VPC specified. `InvalidSecurityGroupIds` returns the list of IDs for security groups that are either invalid or not part of the VPC specified.

HTTP Status Code: 400

InvalidPerUnitStorageThroughput

An invalid value for `PerUnitStorageThroughput` was provided. Please create your file system again, using a valid value.

HTTP Status Code: 400

MissingFileSystemConfiguration

A file system configuration is required for this operation.

HTTP Status Code: 400

ServiceLimitExceeded

An error indicating that a particular service limit was exceeded. You can increase some service limits by contacting AWS Support.

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

DeleteBackup

Deletes an Amazon FSx backup, deleting its contents. After deletion, the backup no longer exists, and its data is gone.

The `DeleteBackup` call returns instantly. The backup will not show up in later `DescribeBackups` calls.

Important

The data in a deleted backup is also deleted and can't be recovered by any means.

Request Syntax

```
{
  "BackupId": "string",
  "ClientRequestToken": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 117\)](#).

The request accepts the following data in JSON format.

BackupId (p. 30)

The ID of the backup you want to delete.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 128.

Pattern: `^(backup-[0-9a-f]{8,})$`

Required: Yes

ClientRequestToken (p. 30)

A string of up to 64 ASCII characters that Amazon FSx uses to ensure idempotent deletion. This is automatically filled on your behalf when using the AWS CLI or SDK.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[A-Za-z0-9_.-]{0,63}$`

Required: No

Response Syntax

```
{
  "BackupId": "string",
  "Lifecycle": "string"
}
```

Response Elements

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

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

BackupId (p. 30)

The ID of the backup deleted.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 128.

Pattern: `^(backup-[0-9a-f]{8,})$`

Lifecycle (p. 30)

The lifecycle of the backup. Should be `DELETED`.

Type: String

Valid Values: `AVAILABLE` | `CREATING` | `TRANSFERRING` | `DELETED` | `FAILED`

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 119).

BackupInProgress

Another backup is already under way. Wait for completion before initiating additional backups of this file system.

HTTP Status Code: 400

BackupNotFound

No Amazon FSx backups were found based upon the supplied parameters.

HTTP Status Code: 400

BackupRestoring

You can't delete a backup while it's being used to restore a file system.

HTTP Status Code: 400

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

IncompatibleParameterError

The error returned when a second request is received with the same client request token but different parameters settings. A client request token should always uniquely identify a single request.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

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

DeleteFileSystem

Deletes a file system, deleting its contents. After deletion, the file system no longer exists, and its data is gone. Any existing automatic backups will also be deleted.

By default, when you delete an Amazon FSx for Windows File Server file system, a final backup is created upon deletion. This final backup is not subject to the file system's retention policy, and must be manually deleted.

The `DeleteFileSystem` action returns while the file system has the `DELETING` status. You can check the file system deletion status by calling the [DescribeFileSystems \(p. 47\)](#) action, which returns a list of file systems in your account. If you pass the file system ID for a deleted file system, the [DescribeFileSystems \(p. 47\)](#) returns a `FileSystemNotFound` error.

Note

Deleting an Amazon FSx for Lustre file system will fail with a 400 `BadRequest` if a data repository task is in a `PENDING` or `EXECUTING` state.

Important

The data in a deleted file system is also deleted and can't be recovered by any means.

Request Syntax

```
{
  "ClientRequestToken": "string",
  "FileSystemId": "string",
  "LustreConfiguration": {
    "FinalBackupTags": [
      {
        "Key": "string",
        "Value": "string"
      }
    ],
    "SkipFinalBackup": boolean
  },
  "WindowsConfiguration": {
    "FinalBackupTags": [
      {
        "Key": "string",
        "Value": "string"
      }
    ],
    "SkipFinalBackup": boolean
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 117\)](#).

The request accepts the following data in JSON format.

ClientRequestToken (p. 33)

A string of up to 64 ASCII characters that Amazon FSx uses to ensure idempotent deletion. This is automatically filled on your behalf when using the AWS CLI or SDK.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[A-Za-z0-9_.-]{0,63}$`

Required: No

FileSystemId (p. 33)

The ID of the file system you want to delete.

Type: String

Length Constraints: Minimum length of 11. Maximum length of 21.

Pattern: `^(fs-[0-9a-f]{8,})$`

Required: Yes

LustreConfiguration (p. 33)

The configuration object for the Amazon FSx for Lustre file system being deleted in the `DeleteFileSystem` operation.

Type: [DeleteFileSystemLustreConfiguration \(p. 90\)](#) object

Required: No

WindowsConfiguration (p. 33)

The configuration object for the Microsoft Windows file system used in the `DeleteFileSystem` operation.

Type: [DeleteFileSystemWindowsConfiguration \(p. 92\)](#) object

Required: No

Response Syntax

```
{
  "FileSystemId": "string",
  "Lifecycle": "string",
  "LustreResponse": {
    "FinalBackupId": "string",
    "FinalBackupTags": [
      {
        "Key": "string",
        "Value": "string"
      }
    ]
  },
  "WindowsResponse": {
    "FinalBackupId": "string",
    "FinalBackupTags": [
      {
        "Key": "string",
        "Value": "string"
      }
    ]
  }
}
```


Response Elements

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

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

FileSystemId (p. 34)

The ID of the file system being deleted.

Type: String

Length Constraints: Minimum length of 11. Maximum length of 21.

Pattern: `^(fs-[0-9a-f]{8,})$`

Lifecycle (p. 34)

The file system lifecycle for the deletion request. Should be `DELETING`.

Type: String

Valid Values: `AVAILABLE` | `CREATING` | `FAILED` | `DELETING` | `MISCONFIGURED` | `UPDATING`

LustreResponse (p. 34)

The response object for the Amazon FSx for Lustre file system being deleted in the `DeleteFileSystem` operation.

Type: [DeleteFileSystemLustreResponse](#) (p. 91) object

WindowsResponse (p. 34)

The response object for the Microsoft Windows file system used in the `DeleteFileSystem` operation.

Type: [DeleteFileSystemWindowsResponse](#) (p. 93) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 119).

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

FileSystemNotFound

No Amazon FSx file systems were found based upon supplied parameters.

HTTP Status Code: 400

IncompatibleParameterError

The error returned when a second request is received with the same client request token but different parameters settings. A client request token should always uniquely identify a single request.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

ServiceLimitExceeded

An error indicating that a particular service limit was exceeded. You can increase some service limits by contacting AWS Support.

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

DescribeBackups

Returns the description of specific Amazon FSx backups, if a `BackupIds` value is provided for that backup. Otherwise, it returns all backups owned by your AWS account in the AWS Region of the endpoint that you're calling.

When retrieving all backups, you can optionally specify the `MaxResults` parameter to limit the number of backups in a response. If more backups remain, Amazon FSx returns a `NextToken` value in the response. In this case, send a later request with the `NextToken` request parameter set to the value of `NextToken` from the last response.

This action is used in an iterative process to retrieve a list of your backups. `DescribeBackups` is called first without a `NextToken` value. Then the action continues to be called with the `NextToken` parameter set to the value of the last `NextToken` value until a response has no `NextToken`.

When using this action, keep the following in mind:

- The implementation might return fewer than `MaxResults` file system descriptions while still including a `NextToken` value.
- The order of backups returned in the response of one `DescribeBackups` call and the order of backups returned across the responses of a multi-call iteration is unspecified.

Request Syntax

```
{
  "BackupIds": [ "string" ],
  "Filters": [
    {
      "Name": "string",
      "Values": [ "string" ]
    }
  ],
  "MaxResults": number,
  "NextToken": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 117\)](#).

The request accepts the following data in JSON format.

BackupIds (p. 37)

IDs of the backups you want to retrieve (String). This overrides any filters. If any IDs are not found, `BackupNotFound` will be thrown.

Type: Array of strings

Array Members: Maximum number of 50 items.

Length Constraints: Minimum length of 12. Maximum length of 128.

Pattern: `^(backup-[0-9a-f]{8,})$`

Required: No

Filters (p. 37)

Filters structure. Supported names are file-system-id and backup-type.

Type: Array of [Filter \(p. 99\)](#) objects

Array Members: Maximum number of 10 items.

Required: No

MaxResults (p. 37)

Maximum number of backups to return in the response (integer). This parameter value must be greater than 0. The number of items that Amazon FSx returns is the minimum of the `MaxResults` parameter specified in the request and the service's internal maximum number of items per page.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 2147483647.

Required: No

NextToken (p. 37)

Opaque pagination token returned from a previous `DescribeBackups` operation (String). If a token present, the action continues the list from where the returning call left off.

Type: String

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

Pattern: `^(?:[A-Za-z0-9+\\/]){4}*(?:[A-Za-z0-9+\\/]){2}==|[A-Za-z0-9+\\/]){3}=?$`

Required: No

Response Syntax

```
{
  "Backups": [
    {
      "BackupId": "string",
      "CreationTime": number,
      "DirectoryInformation": {
        "ActiveDirectoryId": "string",
        "DomainName": "string"
      },
      "FailureDetails": {
        "Message": "string"
      },
      "FileSystem": {
        "AdministrativeActions": [
          {
            "AdministrativeActionType": "string",
            "FailureDetails": {
              "Message": "string"
            },
            "ProgressPercent": number,
            "RequestTime": number,
            "Status": "string",
            "TargetFileSystemValues": "FileSystem"
          }
        ]
      },
      "CreationTime": number,
    }
  ]
}
```

```

    "DNSName": "string",
    "FailureDetails": {
      "Message": "string"
    },
    "FileSystemId": "string",
    "FileSystemType": "string",
    "KmsKeyId": "string",
    "Lifecycle": "string",
    "LustreConfiguration": {
      "AutomaticBackupRetentionDays": number,
      "CopyTagsToBackups": boolean,
      "DailyAutomaticBackupStartTime": "string",
      "DataRepositoryConfiguration": {
        "AutoImportPolicy": "string",
        "ExportPath": "string",
        "FailureDetails": {
          "Message": "string"
        },
        "ImportedFileChunkSize": number,
        "ImportPath": "string",
        "Lifecycle": "string"
      },
      "DeploymentType": "string",
      "DriveCacheType": "string",
      "MountName": "string",
      "PerUnitStorageThroughput": number,
      "WeeklyMaintenanceStartTime": "string"
    },
    "NetworkInterfaceIds": [ "string" ],
    "OwnerId": "string",
    "ResourceARN": "string",
    "StorageCapacity": number,
    "StorageType": "string",
    "SubnetIds": [ "string" ],
    "Tags": [
      {
        "Key": "string",
        "Value": "string"
      }
    ],
    "VpcId": "string",
    "WindowsConfiguration": {
      "ActiveDirectoryId": "string",
      "AutomaticBackupRetentionDays": number,
      "CopyTagsToBackups": boolean,
      "DailyAutomaticBackupStartTime": "string",
      "DeploymentType": "string",
      "MaintenanceOperationsInProgress": [ "string" ],
      "PreferredFileServerIp": "string",
      "PreferredSubnetId": "string",
      "RemoteAdministrationEndpoint": "string",
      "SelfManagedActiveDirectoryConfiguration": {
        "DnsIps": [ "string" ],
        "DomainName": "string",
        "FileSystemAdministratorsGroup": "string",
        "OrganizationalUnitDistinguishedName": "string",
        "UserName": "string"
      },
      "ThroughputCapacity": number,
      "WeeklyMaintenanceStartTime": "string"
    }
  },
  "KmsKeyId": "string",
  "Lifecycle": "string",
  "ProgressPercent": number,
  "ResourceARN": "string",

```

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

Response Elements

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

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

Backups (p. 38)

Any array of backups.

Type: Array of [Backup \(p. 68\)](#) objects

Array Members: Maximum number of 50 items.

NextToken (p. 38)

This is present if there are more backups than returned in the response (String). You can use the `NextToken` value in the later request to fetch the backups.

Type: String

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

Pattern: `^(?:[A-Za-z0-9+\-\/]{4})*(?:[A-Za-z0-9+\-\/]{2}==|[A-Za-z0-9+\-\/]{3}=)?$`

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 119\)](#).

BackupNotFound

No Amazon FSx backups were found based upon the supplied parameters.

HTTP Status Code: 400

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

FileSystemNotFound

No Amazon FSx file systems were found based upon supplied parameters.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

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

DescribeDataRepositoryTasks

Returns the description of specific Amazon FSx for Lustre data repository tasks, if one or more `TaskIds` values are provided in the request, or if filters are used in the request. You can use filters to narrow the response to include just tasks for specific file systems, or tasks in a specific lifecycle state. Otherwise, it returns all data repository tasks owned by your AWS account in the AWS Region of the endpoint that you're calling.

When retrieving all tasks, you can paginate the response by using the optional `MaxResults` parameter to limit the number of tasks returned in a response. If more tasks remain, Amazon FSx returns a `NextToken` value in the response. In this case, send a later request with the `NextToken` request parameter set to the value of `NextToken` from the last response.

Request Syntax

```
{
  "Filters": [
    {
      "Name": "string",
      "Values": [ "string" ]
    }
  ],
  "MaxResults": number,
  "NextToken": "string",
  "TaskIds": [ "string" ]
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 117\)](#).

The request accepts the following data in JSON format.

[Filters \(p. 42\)](#)

(Optional) You can use filters to narrow the `DescribeDataRepositoryTasks` response to include just tasks for specific file systems, or tasks in a specific lifecycle state.

Type: Array of [DataRepositoryTaskFilter \(p. 88\)](#) objects

Array Members: Maximum number of 3 items.

Required: No

[MaxResults \(p. 42\)](#)

The maximum number of resources to return in the response. This value must be an integer greater than zero.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 2147483647.

Required: No

[NextToken \(p. 42\)](#)

(Optional) Opaque pagination token returned from a previous operation (String). If present, this token indicates from what point you can continue processing the request, where the previous `NextToken` value left off.

Type: String

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

Pattern: `^(?:[A-Za-z0-9+\./]{4})*(?:[A-Za-z0-9+\./]{2}==|[A-Za-z0-9+\./]{3}=)?$`

Required: No

TaskIds (p. 42)

(Optional) IDs of the tasks whose descriptions you want to retrieve (String).

Type: Array of strings

Array Members: Maximum number of 50 items.

Length Constraints: Minimum length of 12. Maximum length of 128.

Pattern: `^(task-[0-9a-f]{17,})$`

Required: No

Response Syntax

```
{
  "DataRepositoryTasks": [
    {
      "CreationTime": number,
      "EndTime": number,
      "FailureDetails": {
        "Message": "string"
      },
      "FileSystemId": "string",
      "Lifecycle": "string",
      "Paths": [ "string" ],
      "Report": {
        "Enabled": boolean,
        "Format": "string",
        "Path": "string",
        "Scope": "string"
      },
      "ResourceARN": "string",
      "StartTime": number,
      "Status": {
        "FailedCount": number,
        "LastUpdatedTime": number,
        "SucceededCount": number,
        "TotalCount": number
      },
      "Tags": [
        {
          "Key": "string",
          "Value": "string"
        }
      ],
      "TaskId": "string",
      "Type": "string"
    }
  ],
  "NextToken": "string"
}
```

Response Elements

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

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

DataRepositoryTasks (p. 43)

The collection of data repository task descriptions returned.

Type: Array of [DataRepositoryTask](#) (p. 84) objects

Array Members: Maximum number of 50 items.

NextToken (p. 43)

(Optional) Opaque pagination token returned from a previous operation (String). If present, this token indicates from what point you can continue processing the request, where the previous `NextToken` value left off.

Type: String

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

Pattern: `^(?:[A-Za-z0-9+\-\/]{4})*(?:[A-Za-z0-9+\-\/]{2}==|[A-Za-z0-9+\-\/]{3}=)?$`

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 119).

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

DataRepositoryTaskNotFound

The data repository task or tasks you specified could not be found.

HTTP Status Code: 400

FileSystemNotFound

No Amazon FSx file systems were found based upon supplied parameters.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

Examples

Retrieve Specific Data Repository Task Descriptions

The following request retrieves the descriptions of a specific data repository task by using the `TaskIDs` request parameter.

Sample Request

```
GET /2015-02-01/describe-data-repository-task HTTP/1.1
Host: fsx.us-east-1.amazonaws.com
x-amz-date: 20140620T221118Z
Authorization: <...>
Content-Type: application/json
Content-Length: 160

{
  "TaskIds": ["task-0123456789abcdef0"]
}
```

Sample Response

```
HTTP/1.1 200 success
x-amzn-RequestId: c3616af3-33fa-40ad-ae0d-d3895a2c3a1f

{
  "DataRepositoryTasks": [
    {
      "TaskId": "task-0123456789abcdef0",
      "TaskType": "EXPORT_TO_REPOSITORY",
      "Lifecycle": "PENDING",
      "FileSystemId": "fs-0123456789abcdef1",
      "Paths": ["/path1", "/path2/file1"],
      "CreationTime": "2019-07-17T18:18:18.000Z",
      "TaskReport": {
        "Path": "s3://myBucket/FSxLustre20191118T225838Z/myreports",
        "Format": "REPORT_CSV_20191124",
        "Enabled": true,
        "Scope": "FAILED_FILES_ONLY"
      },
      "Status": {
        "TotalCount": 100,
        "SucceededCount": 0,
        "FailedCount": 0,
        "LastUpdated": "2019-07-17T18:19:05.003Z"
      },
      "Tags": [{"Key": "MyKey"}, {"Value": "MyValue"}],
      "ClientRequestToken": "1234",
      "ResourceARN": "arn:aws:fsx:us-east-1:123456789012:task:task-123f8cd8e330c1321"
    }
  ]
}
```

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

- [AWS SDK for Ruby V3](#)

DescribeFileSystems

Returns the description of specific Amazon FSx file systems, if a `FileSystemIds` value is provided for that file system. Otherwise, it returns descriptions of all file systems owned by your AWS account in the AWS Region of the endpoint that you're calling.

When retrieving all file system descriptions, you can optionally specify the `MaxResults` parameter to limit the number of descriptions in a response. If more file system descriptions remain, Amazon FSx returns a `NextToken` value in the response. In this case, send a later request with the `NextToken` request parameter set to the value of `NextToken` from the last response.

This action is used in an iterative process to retrieve a list of your file system descriptions. `DescribeFileSystems` is called first without a `NextToken` value. Then the action continues to be called with the `NextToken` parameter set to the value of the last `NextToken` value until a response has no `NextToken`.

When using this action, keep the following in mind:

- The implementation might return fewer than `MaxResults` file system descriptions while still including a `NextToken` value.
- The order of file systems returned in the response of one `DescribeFileSystems` call and the order of file systems returned across the responses of a multical iteration is unspecified.

Request Syntax

```
{  
  "FileSystemIds": [ "string" ],  
  "MaxResults": number,  
  "NextToken": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 117\)](#).

The request accepts the following data in JSON format.

[FileSystemIds \(p. 47\)](#)

IDs of the file systems whose descriptions you want to retrieve (String).

Type: Array of strings

Array Members: Maximum number of 50 items.

Length Constraints: Minimum length of 11. Maximum length of 21.

Pattern: `^(fs-[0-9a-f]{8,})$`

Required: No

[MaxResults \(p. 47\)](#)

Maximum number of file systems to return in the response (integer). This parameter value must be greater than 0. The number of items that Amazon FSx returns is the minimum of the `MaxResults` parameter specified in the request and the service's internal maximum number of items per page.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 2147483647.

Required: No

NextToken (p. 47)

Opaque pagination token returned from a previous `DescribeFileSystems` operation (String). If a token present, the action continues the list from where the returning call left off.

Type: String

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

Pattern: `^(?:[A-Za-z0-9+\/]{4})*(?:[A-Za-z0-9+\/]{2}==|[A-Za-z0-9+\/]{3}=)?$`

Required: No

Response Syntax

```
{
  "FileSystems": [
    {
      "AdministrativeActions": [
        {
          "AdministrativeActionType": "string",
          "FailureDetails": {
            "Message": "string"
          },
          "ProgressPercent": number,
          "RequestTime": number,
          "Status": "string",
          "TargetFileSystemValues": "FileSystem"
        }
      ],
      "CreationTime": number,
      "DNSName": "string",
      "FailureDetails": {
        "Message": "string"
      },
      "FileSystemId": "string",
      "FileSystemType": "string",
      "KmsKeyId": "string",
      "Lifecycle": "string",
      "LustreConfiguration": {
        "AutomaticBackupRetentionDays": number,
        "CopyTagsToBackups": boolean,
        "DailyAutomaticBackupStartTime": "string",
        "DataRepositoryConfiguration": {
          "AutoImportPolicy": "string",
          "ExportPath": "string",
          "FailureDetails": {
            "Message": "string"
          },
          "ImportedFileChunkSize": number,
          "ImportPath": "string",
          "Lifecycle": "string"
        },
        "DeploymentType": "string",
        "DriveCacheType": "string",
        "MountName": "string",
        "PerUnitStorageThroughput": number,

```

```

        "WeeklyMaintenanceStartTime": "string"
    },
    "NetworkInterfaceIds": [ "string" ],
    "OwnerId": "string",
    "ResourceARN": "string",
    "StorageCapacity": number,
    "StorageType": "string",
    "SubnetIds": [ "string" ],
    "Tags": [
        {
            "Key": "string",
            "Value": "string"
        }
    ],
    "VpcId": "string",
    "WindowsConfiguration": {
        "ActiveDirectoryId": "string",
        "AutomaticBackupRetentionDays": number,
        "CopyTagsToBackups": boolean,
        "DailyAutomaticBackupStartTime": "string",
        "DeploymentType": "string",
        "MaintenanceOperationsInProgress": [ "string" ],
        "PreferredFileServerIp": "string",
        "PreferredSubnetId": "string",
        "RemoteAdministrationEndpoint": "string",
        "SelfManagedActiveDirectoryConfiguration": {
            "DnsIps": [ "string" ],
            "DomainName": "string",
            "FileSystemAdministratorsGroup": "string",
            "OrganizationalUnitDistinguishedName": "string",
            "UserName": "string"
        },
        "ThroughputCapacity": number,
        "WeeklyMaintenanceStartTime": "string"
    }
}
],
"NextToken": "string"
}

```

Response Elements

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

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

FileSystems (p. 48)

An array of file system descriptions.

Type: Array of [FileSystem \(p. 94\)](#) objects

Array Members: Maximum number of 50 items.

NextToken (p. 48)

Present if there are more file systems than returned in the response (String). You can use the `NextToken` value in the later request to fetch the descriptions.

Type: String

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

Pattern: `^(?:[A-Za-z0-9+\/]{4})*(?:[A-Za-z0-9+\/]{2}==|[A-Za-z0-9+\/]{3}=)?$`

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 119\)](#).

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

FileSystemNotFound

No Amazon FSx file systems were found based upon supplied parameters.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

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

ListTagsForResource

Lists tags for an Amazon FSx file systems and backups in the case of Amazon FSx for Windows File Server.

When retrieving all tags, you can optionally specify the `MaxResults` parameter to limit the number of tags in a response. If more tags remain, Amazon FSx returns a `NextToken` value in the response. In this case, send a later request with the `NextToken` request parameter set to the value of `NextToken` from the last response.

This action is used in an iterative process to retrieve a list of your tags. `ListTagsForResource` is called first without a `NextToken` value. Then the action continues to be called with the `NextToken` parameter set to the value of the last `NextToken` value until a response has no `NextToken`.

When using this action, keep the following in mind:

- The implementation might return fewer than `MaxResults` file system descriptions while still including a `NextToken` value.
- The order of tags returned in the response of one `ListTagsForResource` call and the order of tags returned across the responses of a multi-call iteration is unspecified.

Request Syntax

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

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 117\)](#).

The request accepts the following data in JSON format.

MaxResults (p. 51)

Maximum number of tags to return in the response (integer). This parameter value must be greater than 0. The number of items that Amazon FSx returns is the minimum of the `MaxResults` parameter specified in the request and the service's internal maximum number of items per page.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 2147483647.

Required: No

NextToken (p. 51)

Opaque pagination token returned from a previous `ListTagsForResource` operation (String). If a token present, the action continues the list from where the returning call left off.

Type: String

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

Pattern: `^(?:[A-Za-z0-9+\/]{4})*(?:[A-Za-z0-9+\/]{2}==|[A-Za-z0-9+\/]{3}=)?$`

Required: No

ResourceARN (p. 51)

The ARN of the Amazon FSx resource that will have its tags listed.

Type: String

Length Constraints: Minimum length of 8. Maximum length of 512.

Pattern: `^arn:(?=[^:]+:fsx:[^:]+\d{12}:)(\(|(?=[a-z0-9-.]){1,63})(?!\\d{1,3})(\\.\\d{1,3}){3})(?![^\:]*-{2})(?![^\:]*-\\.)(?![^\:]*\\.-)[a-z0-9].*(?!-))){4}(?!/).\{0,1024}$`

Required: Yes

Response Syntax

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

Response Elements

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

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

NextToken (p. 52)

This is present if there are more tags than returned in the response (String). You can use the `NextToken` value in the later request to fetch the tags.

Type: String

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

Pattern: `^(?:[A-Za-z0-9+\/]{4})*(?:[A-Za-z0-9+\/]{2}==|[A-Za-z0-9+\/]{3}=)?$`

Tags (p. 52)

A list of tags on the resource.

Type: Array of [Tag \(p. 109\)](#) objects

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

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 119\)](#).

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

NotServiceResourceError

The resource specified for the tagging operation is not a resource type owned by Amazon FSx. Use the API of the relevant service to perform the operation.

HTTP Status Code: 400

ResourceDoesNotSupportTagging

The resource specified does not support tagging.

HTTP Status Code: 400

ResourceNotFound

The resource specified by the Amazon Resource Name (ARN) can't be found.

HTTP Status Code: 400

See Also

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

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

TagResource

Tags an Amazon FSx resource.

Request Syntax

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

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 117\)](#).

The request accepts the following data in JSON format.

ResourceARN (p. 54)

The Amazon Resource Name (ARN) of the Amazon FSx resource that you want to tag.

Type: String

Length Constraints: Minimum length of 8. Maximum length of 512.

Pattern: `^arn:(?=[^:]+:fsx:[^:]+\d{12}:)(\(|(?=[a-z0-9-.]{1,63})(?!\\d{1,3}(\.\\d{1,3}){3})(?![^\:]*-){2})(?![^\:]*-\.)(?![^\:]*\.-)[a-z0-9].*(?!-)):\d{4}(?!/)\. {0,1024}$`

Required: Yes

Tags (p. 54)

A list of tags for the resource. If a tag with a given key already exists, the value is replaced by the one specified in this parameter.

Type: Array of [Tag \(p. 109\)](#) objects

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

Required: Yes

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 119\)](#).

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

NotServiceResourceError

The resource specified for the tagging operation is not a resource type owned by Amazon FSx. Use the API of the relevant service to perform the operation.

HTTP Status Code: 400

ResourceDoesNotSupportTagging

The resource specified does not support tagging.

HTTP Status Code: 400

ResourceNotFound

The resource specified by the Amazon Resource Name (ARN) can't be found.

HTTP Status Code: 400

See Also

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

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

UntagResource

This action removes a tag from an Amazon FSx resource.

Request Syntax

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

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 117\)](#).

The request accepts the following data in JSON format.

ResourceARN (p. 56)

The ARN of the Amazon FSx resource to untag.

Type: String

Length Constraints: Minimum length of 8. Maximum length of 512.

Pattern: `^arn:(?=[^:]+:fsx:[^:]+\d{12}:)(\(|(?=[a-z0-9-.\]{1,63})(?!\\d{1,3}(\.\\d{1,3}){3})(?![^\:]*-{2})(?![^\:]*-\.)(?![^\:]*\.-)[a-z0-9].*(?!-)):\d{4}(?!/)\. {0,1024}$`

Required: Yes

TagKeys (p. 56)

A list of keys of tags on the resource to untag. In case the tag key doesn't exist, the call will still succeed to be idempotent.

Type: Array of strings

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

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

Pattern: `^[\\p{L}\\p{Z}\\p{N}_.:/=+\\-@]*$`

Required: Yes

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 119\)](#).

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

NotServiceResourceError

The resource specified for the tagging operation is not a resource type owned by Amazon FSx. Use the API of the relevant service to perform the operation.

HTTP Status Code: 400

ResourceDoesNotSupportTagging

The resource specified does not support tagging.

HTTP Status Code: 400

ResourceNotFound

The resource specified by the Amazon Resource Name (ARN) can't be found.

HTTP Status Code: 400

See Also

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

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

UpdateFileSystem

Use this operation to update the configuration of an existing Amazon FSx file system. You can update multiple properties in a single request.

For Amazon FSx for Windows File Server file systems, you can update the following properties:

- AutomaticBackupRetentionDays
- DailyAutomaticBackupStartTime
- SelfManagedActiveDirectoryConfiguration
- StorageCapacity
- ThroughputCapacity
- WeeklyMaintenanceStartTime

For Amazon FSx for Lustre file systems, you can update the following properties:

- AutoImportPolicy
- AutomaticBackupRetentionDays
- DailyAutomaticBackupStartTime
- WeeklyMaintenanceStartTime

Request Syntax

```
{
  "ClientRequestToken": "string",
  "FileSystemId": "string",
  "LustreConfiguration": {
    "AutoImportPolicy": "string",
    "AutomaticBackupRetentionDays": number,
    "DailyAutomaticBackupStartTime": "string",
    "WeeklyMaintenanceStartTime": "string"
  },
  "StorageCapacity": number,
  "WindowsConfiguration": {
    "AutomaticBackupRetentionDays": number,
    "DailyAutomaticBackupStartTime": "string",
    "SelfManagedActiveDirectoryConfiguration": {
      "DnsIps": [ "string" ],
      "Password": "string",
      "UserName": "string"
    },
    "ThroughputCapacity": number,
    "WeeklyMaintenanceStartTime": "string"
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 117).

The request accepts the following data in JSON format.

ClientRequestToken (p. 58)

A string of up to 64 ASCII characters that Amazon FSx uses to ensure idempotent updates. This string is automatically filled on your behalf when you use the AWS Command Line Interface (AWS CLI) or an AWS SDK.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[A-Za-z0-9_.-]{0,63}`\$

Required: No

FileSystemId (p. 58)

Identifies the file system that you are updating.

Type: String

Length Constraints: Minimum length of 11. Maximum length of 21.

Pattern: `^(fs-[0-9a-f]{8,})$`

Required: Yes

LustreConfiguration (p. 58)

The configuration object for Amazon FSx for Lustre file systems used in the `UpdateFileSystem` operation.

Type: [UpdateFileSystemLustreConfiguration \(p. 110\)](#) object

Required: No

StorageCapacity (p. 58)

Use this parameter to increase the storage capacity of an Amazon FSx for Windows File Server file system. Specifies the storage capacity target value, GiB, for the file system you're updating. The storage capacity target value must be at least 10 percent (%) greater than the current storage capacity value. In order to increase storage capacity, the file system needs to have at least 16 MB/s of throughput capacity. You cannot make a storage capacity increase request if there is an existing storage capacity increase request in progress. For more information, see [Managing Storage Capacity](#).

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 2147483647.

Required: No

WindowsConfiguration (p. 58)

The configuration updates for an Amazon FSx for Windows File Server file system.

Type: [UpdateFileSystemWindowsConfiguration \(p. 112\)](#) object

Required: No

Response Syntax

```
{  
  "FileSystem": {
```

```

"AdministrativeActions": [
  {
    "AdministrativeActionType": "string",
    "FailureDetails": {
      "Message": "string"
    },
    "ProgressPercent": number,
    "RequestTime": number,
    "Status": "string",
    "TargetFileSystemValues": "FileSystem"
  }
],
"CreationTime": number,
"DNSName": "string",
"FailureDetails": {
  "Message": "string"
},
"FileSystemId": "string",
"FileSystemType": "string",
"KmsKeyId": "string",
"Lifecycle": "string",
"LustreConfiguration": {
  "AutomaticBackupRetentionDays": number,
  "CopyTagsToBackups": boolean,
  "DailyAutomaticBackupStartTime": "string",
  "DataRepositoryConfiguration": {
    "AutoImportPolicy": "string",
    "ExportPath": "string",
    "FailureDetails": {
      "Message": "string"
    },
    "ImportedFileChunkSize": number,
    "ImportPath": "string",
    "Lifecycle": "string"
  },
  "DeploymentType": "string",
  "DriveCacheType": "string",
  "MountName": "string",
  "PerUnitStorageThroughput": number,
  "WeeklyMaintenanceStartTime": "string"
},
"NetworkInterfaceIds": [ "string" ],
"OwnerId": "string",
"ResourceARN": "string",
"StorageCapacity": number,
"StorageType": "string",
"SubnetIds": [ "string" ],
"Tags": [
  {
    "Key": "string",
    "Value": "string"
  }
],
"VpcId": "string",
"WindowsConfiguration": {
  "ActiveDirectoryId": "string",
  "AutomaticBackupRetentionDays": number,
  "CopyTagsToBackups": boolean,
  "DailyAutomaticBackupStartTime": "string",
  "DeploymentType": "string",
  "MaintenanceOperationsInProgress": [ "string" ],
  "PreferredFileServerIp": "string",
  "PreferredSubnetId": "string",
  "RemoteAdministrationEndpoint": "string",
  "SelfManagedActiveDirectoryConfiguration": {
    "DnsIps": [ "string" ],

```

```
    "DomainName": "string",  
    "FileSystemAdministratorsGroup": "string",  
    "OrganizationalUnitDistinguishedName": "string",  
    "UserName": "string"  
  },  
  "ThroughputCapacity": number,  
  "WeeklyMaintenanceStartTime": "string"  
}  
}
```

Response Elements

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

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

FileSystem (p. 59)

A description of the file system that was updated.

Type: [FileSystem \(p. 94\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 119\)](#).

BadRequest

A generic error indicating a failure with a client request.

HTTP Status Code: 400

FileSystemNotFound

No Amazon FSx file systems were found based upon supplied parameters.

HTTP Status Code: 400

IncompatibleParameterError

The error returned when a second request is received with the same client request token but different parameters settings. A client request token should always uniquely identify a single request.

HTTP Status Code: 400

InternalServerError

A generic error indicating a server-side failure.

HTTP Status Code: 500

MissingFileSystemConfiguration

A file system configuration is required for this operation.

HTTP Status Code: 400

ServiceLimitExceeded

An error indicating that a particular service limit was exceeded. You can increase some service limits by contacting AWS Support.

HTTP Status Code: 400

UnsupportedOperation

The requested operation is not supported for this resource or API.

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

Data Types

The Amazon FSx 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:

- [ActiveDirectoryBackupAttributes](#) (p. 64)
- [AdministrativeAction](#) (p. 65)
- [AdministrativeActionFailureDetails](#) (p. 67)
- [Backup](#) (p. 68)
- [BackupFailureDetails](#) (p. 71)
- [CompletionReport](#) (p. 72)
- [CreateFileSystemLustreConfiguration](#) (p. 74)
- [CreateFileSystemWindowsConfiguration](#) (p. 78)
- [DataRepositoryConfiguration](#) (p. 81)
- [DataRepositoryFailureDetails](#) (p. 83)
- [DataRepositoryTask](#) (p. 84)
- [DataRepositoryTaskFailureDetails](#) (p. 87)
- [DataRepositoryTaskFilter](#) (p. 88)
- [DataRepositoryTaskStatus](#) (p. 89)
- [DeleteFileSystemLustreConfiguration](#) (p. 90)
- [DeleteFileSystemLustreResponse](#) (p. 91)
- [DeleteFileSystemWindowsConfiguration](#) (p. 92)
- [DeleteFileSystemWindowsResponse](#) (p. 93)
- [FileSystem](#) (p. 94)
- [FileSystemFailureDetails](#) (p. 98)
- [Filter](#) (p. 99)
- [LustreFileSystemConfiguration](#) (p. 100)
- [SelfManagedActiveDirectoryAttributes](#) (p. 103)
- [SelfManagedActiveDirectoryConfiguration](#) (p. 105)
- [SelfManagedActiveDirectoryConfigurationUpdates](#) (p. 107)
- [Tag](#) (p. 109)
- [UpdateFileSystemLustreConfiguration](#) (p. 110)
- [UpdateFileSystemWindowsConfiguration](#) (p. 112)
- [WindowsFileSystemConfiguration](#) (p. 114)

ActiveDirectoryBackupAttributes

The Microsoft AD attributes of the Amazon FSx for Windows File Server file system.

Contents

ActiveDirectoryId

The ID of the AWS Managed Microsoft Active Directory instance to which the file system is joined.

Type: String

Length Constraints: Fixed length of 12.

Pattern: `^d-[0-9a-f]{10}$`

Required: No

DomainName

The fully qualified domain name of the self-managed AD directory.

Type: String

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

Pattern: `^[^\u0000\u0085\u2028\u2029\r\n]{1,255}$`

Required: No

See Also

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

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

AdministrativeAction

Describes a specific Amazon FSx Administrative Action for the current Windows file system.

Contents

AdministrativeActionType

Describes the type of administrative action, as follows:

- `FILE_SYSTEM_UPDATE` - A file system update administrative action initiated by the user from the Amazon FSx console, API (`UpdateFileSystem`), or CLI (`update-file-system`). A
- `STORAGE_OPTIMIZATION` - Once the `FILE_SYSTEM_UPDATE` task to increase a file system's storage capacity completes successfully, a `STORAGE_OPTIMIZATION` task starts. Storage optimization is the process of migrating the file system data to the new, larger disks. You can track the storage migration progress using the `ProgressPercent` property. When `STORAGE_OPTIMIZATION` completes successfully, the parent `FILE_SYSTEM_UPDATE` action status changes to `COMPLETED`. For more information, see [Managing Storage Capacity](#).

Type: String

Valid Values: `FILE_SYSTEM_UPDATE` | `STORAGE_OPTIMIZATION`

Required: No

FailureDetails

Provides information about a failed administrative action.

Type: [AdministrativeActionFailureDetails](#) (p. 67) object

Required: No

ProgressPercent

Provides the percent complete of a `STORAGE_OPTIMIZATION` administrative action.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 100.

Required: No

RequestTime

Time that the administrative action request was received.

Type: Timestamp

Required: No

Status

Describes the status of the administrative action, as follows:

- `FAILED` - Amazon FSx failed to process the administrative action successfully.
- `IN_PROGRESS` - Amazon FSx is processing the administrative action.
- `PENDING` - Amazon FSx is waiting to process the administrative action.
- `COMPLETED` - Amazon FSx has finished processing the administrative task.
- `UPDATED_OPTIMIZING` - For a storage capacity increase update, Amazon FSx has updated the file system with the new storage capacity, and is now performing the storage optimization process. For more information, see [Managing Storage Capacity](#).

Type: String

Valid Values: FAILED | IN_PROGRESS | PENDING | COMPLETED | UPDATED_OPTIMIZING

Required: No

TargetFileSystemValues

Describes the target `StorageCapacity` or `ThroughputCapacity` value provided in the `UpdateFileSystem` operation. Returned for `FILE_SYSTEM_UPDATE` administrative actions.

Type: [FileSystem \(p. 94\)](#) object

Required: No

See Also

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

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

AdministrativeActionFailureDetails

Provides information about a failed administrative action.

Contents

Message

Error message providing details about the failure.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

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

Backup

A backup of an Amazon FSx for file system.

Contents

BackupId

The ID of the backup.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 128.

Pattern: `^(backup-[0-9a-f]{8,})$`

Required: Yes

CreationTime

The time when a particular backup was created.

Type: Timestamp

Required: Yes

DirectoryInformation

The configuration of the self-managed Microsoft Active Directory (AD) to which the Windows File Server instance is joined.

Type: [ActiveDirectoryBackupAttributes \(p. 64\)](#) object

Required: No

FailureDetails

Details explaining any failures that occur when creating a backup.

Type: [BackupFailureDetails \(p. 71\)](#) object

Required: No

FileSystem

Metadata of the file system associated with the backup. This metadata is persisted even if the file system is deleted.

Type: [FileSystem \(p. 94\)](#) object

Required: Yes

KmsKeyId

The ID of the AWS Key Management Service (AWS KMS) key used to encrypt the backup of the Amazon FSx file system's data at rest.

Type: String

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

Pattern: `^.{1,2048}$`

Required: No

Lifecycle

The lifecycle status of the backup.

- **AVAILABLE** - The backup is fully available.
- **CREATING** - FSx is creating the backup.
- **TRANSFERRING** - For Lustre file systems only; FSx is transferring the backup to S3.
- **DELETED** - The backup was deleted is no longer available.
- **FAILED** - Amazon FSx could not complete the backup.

Type: String

Valid Values: **AVAILABLE** | **CREATING** | **TRANSFERRING** | **DELETED** | **FAILED**

Required: Yes

ProgressPercent

The current percent of progress of an asynchronous task.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 100.

Required: No

ResourceARN

The Amazon Resource Name (ARN) for the backup resource.

Type: String

Length Constraints: Minimum length of 8. Maximum length of 512.

Pattern: `^arn:(?=[^:]+:fsx:[^:]+\d{12}:)(\(|(?=[a-z0-9-]{1,63})(?!\\d{1,3}|\.\\d{1,3}){3})(?![^\s]*-){2})(?![^\s]*-\.)(?![^\s]*\\\.)[a-z0-9].*(?!-))){4}(?!/)\. {0,1024}$`

Required: No

Tags

Tags associated with a particular file system.

Type: Array of [Tag \(p. 109\)](#) objects

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

Required: No

Type

The type of the file system backup.

Type: String

Valid Values: **AUTOMATIC** | **USER_INITIATED**

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

BackupFailureDetails

If backup creation fails, this structure contains the details of that failure.

Contents

Message

A message describing the backup creation failure.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

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

CompletionReport

Provides a report detailing the data repository task results of the files processed that match the criteria specified in the report `Scope` parameter. FSx delivers the report to the file system's linked data repository in Amazon S3, using the path specified in the report `Path` parameter. You can specify whether or not a report gets generated for a task using the `Enabled` parameter.

Contents

Enabled

Set `Enabled` to `True` to generate a `CompletionReport` when the task completes. If set to `true`, then you need to provide a report `Scope`, `Path`, and `Format`. Set `Enabled` to `False` if you do not want a `CompletionReport` generated when the task completes.

Type: Boolean

Required: Yes

Format

Required if `Enabled` is set to `true`. Specifies the format of the `CompletionReport`. `REPORT_CSV_20191124` is the only format currently supported. When `Format` is set to `REPORT_CSV_20191124`, the `CompletionReport` is provided in CSV format, and is delivered to `{path}/task-{id}/failures.csv`.

Type: String

Valid Values: `REPORT_CSV_20191124`

Required: No

Path

Required if `Enabled` is set to `true`. Specifies the location of the report on the file system's linked S3 data repository. An absolute path that defines where the completion report will be stored in the destination location. The `Path` you provide must be located within the file system's `ExportPath`. An example `Path` value is `"s3://myBucket/myExportPath/optionalPrefix"`. The report provides the following information for each file in the report: `FilePath`, `FileStatus`, and `ErrorCode`. To learn more about a file system's `ExportPath`, see [DataRepositoryConfiguration \(p. 81\)](#).

Type: String

Length Constraints: Minimum length of 3. Maximum length of 900.

Pattern: `^[^\u0000\u0085\u2028\u2029\r\n]{3,4357}$`

Required: No

Scope

Required if `Enabled` is set to `true`. Specifies the scope of the `CompletionReport`; `FAILED_FILES_ONLY` is the only scope currently supported. When `Scope` is set to `FAILED_FILES_ONLY`, the `CompletionReport` only contains information about files that the data repository task failed to process.

Type: String

Valid Values: `FAILED_FILES_ONLY`

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

CreateFileSystemLustreConfiguration

The Lustre configuration for the file system being created.

Contents

AutoImportPolicy

(Optional) When you create your file system, your existing S3 objects appear as file and directory listings. Use this property to choose how Amazon FSx keeps your file and directory listings up to date as you add or modify objects in your linked S3 bucket. `AutoImportPolicy` can have the following values:

- `NONE` - (Default) AutoImport is off. Amazon FSx only updates file and directory listings from the linked S3 bucket when the file system is created. FSx does not update file and directory listings for any new or changed objects after choosing this option.
- `NEW` - AutoImport is on. Amazon FSx automatically imports directory listings of any new objects added to the linked S3 bucket that do not currently exist in the FSx file system.
- `NEW_CHANGED` - AutoImport is on. Amazon FSx automatically imports file and directory listings of any new objects added to the S3 bucket and any existing objects that are changed in the S3 bucket after you choose this option.

For more information, see [Automatically import updates from your S3 bucket](#).

Type: String

Valid Values: `NONE` | `NEW` | `NEW_CHANGED`

Required: No

AutomaticBackupRetentionDays

The number of days to retain automatic backups. Setting this to 0 disables automatic backups. You can retain automatic backups for a maximum of 90 days. The default is 0.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 90.

Required: No

CopyTagsToBackups

(Optional) Not available to use with file systems that are linked to a data repository. A boolean flag indicating whether tags for the file system should be copied to backups. The default value is false. If it's set to true, all file system tags are copied to all automatic and user-initiated backups when the user doesn't specify any backup-specific tags. If this value is true, and you specify one or more backup tags, only the specified tags are copied to backups. If you specify one or more tags when creating a user-initiated backup, no tags are copied from the file system, regardless of this value.

For more information, see [Working with backups](#).

Type: Boolean

Required: No

DailyAutomaticBackupStartTime

A recurring daily time, in the format `HH:MM`. `HH` is the zero-padded hour of the day (0-23), and `MM` is the zero-padded minute of the hour. For example, `05:00` specifies 5 AM daily.

Type: String

Length Constraints: Fixed length of 5.

Pattern: `^([01]\d|2[0-3]):?([0-5]\d)$`

Required: No

DeploymentType

Choose `SCRATCH_1` and `SCRATCH_2` deployment types when you need temporary storage and shorter-term processing of data. The `SCRATCH_2` deployment type provides in-transit encryption of data and higher burst throughput capacity than `SCRATCH_1`.

Choose `PERSISTENT_1` deployment type for longer-term storage and workloads and encryption of data in transit. To learn more about deployment types, see [FSx for Lustre Deployment Options](#).

Encryption of data in-transit is automatically enabled when you access a `SCRATCH_2` or `PERSISTENT_1` file system from Amazon EC2 instances that [support this feature](#). (Default = `SCRATCH_1`)

Encryption of data in-transit for `SCRATCH_2` and `PERSISTENT_1` deployment types is supported when accessed from supported instance types in supported AWS Regions. To learn more, [Encrypting Data in Transit](#).

Type: String

Valid Values: `SCRATCH_1` | `SCRATCH_2` | `PERSISTENT_1`

Required: No

DriveCacheType

The type of drive cache used by `PERSISTENT_1` file systems that are provisioned with HDD storage devices. This parameter is required when storage type is HDD. Set to `READ`, improve the performance for frequently accessed files and allows 20% of the total storage capacity of the file system to be cached.

This parameter is required when `StorageType` is set to HDD.

Type: String

Valid Values: `NONE` | `READ`

Required: No

ExportPath

(Optional) The path in Amazon S3 where the root of your Amazon FSx file system is exported. The path must use the same Amazon S3 bucket as specified in `ImportPath`. You can provide an optional prefix to which new and changed data is to be exported from your Amazon FSx for Lustre file system. If an `ExportPath` value is not provided, Amazon FSx sets a default export path, `s3://import-bucket/FSxLustre[creation-timestamp]`. The timestamp is in UTC format, for example `s3://import-bucket/FSxLustre20181105T222312Z`.

The Amazon S3 export bucket must be the same as the import bucket specified by `ImportPath`. If you only specify a bucket name, such as `s3://import-bucket`, you get a 1:1 mapping of file system objects to S3 bucket objects. This mapping means that the input data in S3 is overwritten on export. If you provide a custom prefix in the export path, such as `s3://import-bucket/[custom-optional-prefix]`, Amazon FSx exports the contents of your file system to that export prefix in the Amazon S3 bucket.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 900.

Pattern: `^[^\u0000\u0085\u2028\u2029\r\n]{3,4357}$`

Required: No

ImportedFileChunkSize

(Optional) For files imported from a data repository, this value determines the stripe count and maximum amount of data per file (in MiB) stored on a single physical disk. The maximum number of disks that a single file can be striped across is limited by the total number of disks that make up the file system.

The default chunk size is 1,024 MiB (1 GiB) and can go as high as 512,000 MiB (500 GiB). Amazon S3 objects have a maximum size of 5 TB.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 512000.

Required: No

ImportPath

(Optional) The path to the Amazon S3 bucket (including the optional prefix) that you're using as the data repository for your Amazon FSx for Lustre file system. The root of your FSx for Lustre file system will be mapped to the root of the Amazon S3 bucket you select. An example is `s3://import-bucket/optional-prefix`. If you specify a prefix after the Amazon S3 bucket name, only object keys with that prefix are loaded into the file system.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 900.

Pattern: `^[^\u0000\u0085\u2028\u2029\r\n]{3,4357}$`

Required: No

PerUnitStorageThroughput

Required for the `PERSISTENT_1` deployment type, describes the amount of read and write throughput for each 1 tebibyte of storage, in MB/s/TiB. File system throughput capacity is calculated by multiplying file system storage capacity (TiB) by the `PerUnitStorageThroughput` (MB/s/TiB). For a 2.4 TiB file system, provisioning 50 MB/s/TiB of `PerUnitStorageThroughput` yields 120 MB/s of file system throughput. You pay for the amount of throughput that you provision.

Valid values for SSD storage: 50, 100, 200. Valid values for HDD storage: 12, 40.

Type: Integer

Valid Range: Minimum value of 12. Maximum value of 200.

Required: No

WeeklyMaintenanceStartTime

(Optional) The preferred start time to perform weekly maintenance, formatted d:HH:MM in the UTC time zone, where d is the weekday number, from 1 through 7, beginning with Monday and ending with Sunday.

Type: String

Length Constraints: Fixed length of 7.

Pattern: `^[1-7]:([01]\d|2[0-3]):?([0-5]\d)$`

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

CreateFileSystemWindowsConfiguration

The configuration object for the Microsoft Windows file system used in `CreateFileSystem` and `CreateFileSystemFromBackup` operations.

Contents

ActiveDirectoryId

The ID for an existing AWS Managed Microsoft Active Directory (AD) instance that the file system should join when it's created.

Type: String

Length Constraints: Fixed length of 12.

Pattern: `^d-[0-9a-f]{10}$`

Required: No

AutomaticBackupRetentionDays

The number of days to retain automatic backups. The default is to retain backups for 7 days. Setting this value to 0 disables the creation of automatic backups. The maximum retention period for backups is 90 days.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 90.

Required: No

CopyTagsToBackups

A boolean flag indicating whether tags for the file system should be copied to backups. This value defaults to false. If it's set to true, all tags for the file system are copied to all automatic and user-initiated backups where the user doesn't specify tags. If this value is true, and you specify one or more tags, only the specified tags are copied to backups. If you specify one or more tags when creating a user-initiated backup, no tags are copied from the file system, regardless of this value.

Type: Boolean

Required: No

DailyAutomaticBackupStartTime

The preferred time to take daily automatic backups, formatted HH:MM in the UTC time zone.

Type: String

Length Constraints: Fixed length of 5.

Pattern: `^([01]\d|2[0-3]):?([0-5]\d)$`

Required: No

DeploymentType

Specifies the file system deployment type, valid values are the following:

- `MULTI_AZ_1` - Deploys a high availability file system that is configured for Multi-AZ redundancy to tolerate temporary Availability Zone (AZ) unavailability. You can only deploy a Multi-AZ file

system in AWS Regions that have a minimum of three Availability Zones. Also supports HDD storage type

- `SINGLE_AZ_1` - (Default) Choose to deploy a file system that is configured for single AZ redundancy.
- `SINGLE_AZ_2` - The latest generation Single AZ file system. Specifies a file system that is configured for single AZ redundancy and supports HDD storage type.

For more information, see [Availability and Durability: Single-AZ and Multi-AZ File Systems](#).

Type: String

Valid Values: `MULTI_AZ_1` | `SINGLE_AZ_1` | `SINGLE_AZ_2`

Required: No

PreferredSubnetId

Required when `DeploymentType` is set to `MULTI_AZ_1`. This specifies the subnet in which you want the preferred file server to be located. For in-AWS applications, we recommend that you launch your clients in the same Availability Zone (AZ) as your preferred file server to reduce cross-AZ data transfer costs and minimize latency.

Type: String

Length Constraints: Minimum length of 15. Maximum length of 24.

Pattern: `^(subnet-[0-9a-f]{8,})$`

Required: No

SelfManagedActiveDirectoryConfiguration

The configuration that Amazon FSx uses to join the Windows File Server instance to your self-managed (including on-premises) Microsoft Active Directory (AD) directory.

Type: [SelfManagedActiveDirectoryConfiguration \(p. 105\)](#) object

Required: No

ThroughputCapacity

The throughput of an Amazon FSx file system, measured in megabytes per second, in 2 to the *n*th increments, between 2³ (8) and 2¹¹ (2048).

Type: Integer

Valid Range: Minimum value of 8. Maximum value of 2048.

Required: Yes

WeeklyMaintenanceStartTime

The preferred start time to perform weekly maintenance, formatted d:HH:MM in the UTC time zone, where d is the weekday number, from 1 through 7, beginning with Monday and ending with Sunday.

Type: String

Length Constraints: Fixed length of 7.

Pattern: `^[1-7]:([01]\d|2[0-3]):?([0-5]\d)$`

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

DataRepositoryConfiguration

The data repository configuration object for Lustre file systems returned in the response of the `CreateFileSystem` operation.

Contents

AutoImportPolicy

Describes the file system's linked S3 data repository's `AutoImportPolicy`. The `AutoImportPolicy` configures how Amazon FSx keeps your file and directory listings up to date as you add or modify objects in your linked S3 bucket. `AutoImportPolicy` can have the following values:

- `NONE` - (Default) `AutoImport` is off. Amazon FSx only updates file and directory listings from the linked S3 bucket when the file system is created. FSx does not update file and directory listings for any new or changed objects after choosing this option.
- `NEW` - `AutoImport` is on. Amazon FSx automatically imports directory listings of any new objects added to the linked S3 bucket that do not currently exist in the FSx file system.
- `NEW_CHANGED` - `AutoImport` is on. Amazon FSx automatically imports file and directory listings of any new objects added to the S3 bucket and any existing objects that are changed in the S3 bucket after you choose this option.

For more information, see [Automatically import updates from your S3 bucket](#).

Type: String

Valid Values: `NONE` | `NEW` | `NEW_CHANGED`

Required: No

ExportPath

The export path to the Amazon S3 bucket (and prefix) that you are using to store new and changed Lustre file system files in S3.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 900.

Pattern: `^[^\u0000\u0085\u2028\u2029\r\n]{3,4357}$`

Required: No

FailureDetails

Provides detailed information about the data repository if its `Lifecycle` is set to `MISCONFIGURED`.

Type: [DataRepositoryFailureDetails](#) (p. 83) object

Required: No

ImportedFileChunkSize

For files imported from a data repository, this value determines the stripe count and maximum amount of data per file (in MiB) stored on a single physical disk. The maximum number of disks that a single file can be striped across is limited by the total number of disks that make up the file system.

The default chunk size is 1,024 MiB (1 GiB) and can go as high as 512,000 MiB (500 GiB). Amazon S3 objects have a maximum size of 5 TB.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 512000.

Required: No

ImportPath

The import path to the Amazon S3 bucket (and optional prefix) that you're using as the data repository for your FSx for Lustre file system, for example `s3://import-bucket/optional-prefix`. If a prefix is specified after the Amazon S3 bucket name, only object keys with that prefix are loaded into the file system.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 900.

Pattern: `^[^\u0000\u0085\u2028\u2029\r\n]{3,4357}$`

Required: No

Lifecycle

Describes the state of the file system's S3 durable data repository, if it is configured with an S3 repository. The lifecycle can have the following values:

- **CREATING** - The data repository configuration between the FSx file system and the linked S3 data repository is being created. The data repository is unavailable.
- **AVAILABLE** - The data repository is available for use.
- **MISCONFIGURED** - Amazon FSx cannot automatically import updates from the S3 bucket until the data repository configuration is corrected. For more information, see [Troubleshooting a Misconfigured linked S3 bucket](#).
- **UPDATING** - The data repository is undergoing a customer initiated update and availability may be impacted.

Type: String

Valid Values: `CREATING | AVAILABLE | MISCONFIGURED | UPDATING | DELETING`

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

DataRepositoryFailureDetails

Provides detailed information about the data repository if its `Lifecycle` is set to `MISCONFIGURED`.

Contents

Message

A detailed error message.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

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

DataRepositoryTask

A description of the data repository task. You use data repository tasks to perform bulk transfer operations between your Amazon FSx file system and its linked data repository.

Contents

CreationTime

The time that the resource was created, in seconds (since 1970-01-01T00:00:00Z), also known as Unix time.

Type: Timestamp

Required: Yes

EndTime

The time that Amazon FSx completed processing the task, populated after the task is complete.

Type: Timestamp

Required: No

FailureDetails

Failure message describing why the task failed, it is populated only when `Lifecycle` is set to `FAILED`.

Type: [DataRepositoryTaskFailureDetails \(p. 87\)](#) object

Required: No

FileSystemId

The globally unique ID of the file system, assigned by Amazon FSx.

Type: String

Length Constraints: Minimum length of 11. Maximum length of 21.

Pattern: `^(fs-[0-9a-f]{8,})$`

Required: Yes

Lifecycle

The lifecycle status of the data repository task, as follows:

- `PENDING` - Amazon FSx has not started the task.
- `EXECUTING` - Amazon FSx is processing the task.
- `FAILED` - Amazon FSx was not able to complete the task. For example, there may be files the task failed to process. The [DataRepositoryTaskFailureDetails \(p. 87\)](#) property provides more information about task failures.
- `SUCCEEDED` - FSx completed the task successfully.
- `CANCELED` - Amazon FSx canceled the task and it did not complete.
- `CANCELING` - FSx is in process of canceling the task.

Note

You cannot delete an FSx for Lustre file system if there are data repository tasks for the file system in the `PENDING` or `EXECUTING` states. Please retry when the data repository

task is finished (with a status of CANCELED, SUCCEEDED, or FAILED). You can use the DescribeDataRepositoryTask action to monitor the task status. Contact the FSx team if you need to delete your file system immediately.

Type: String

Valid Values: PENDING | EXECUTING | FAILED | SUCCEEDED | CANCELED | CANCELING

Required: Yes

Paths

An array of paths on the Amazon FSx for Lustre file system that specify the data for the data repository task to process. For example, in an EXPORT_TO_REPOSITORY task, the paths specify which data to export to the linked data repository.

(Default) If Paths is not specified, Amazon FSx uses the file system root directory.

Type: Array of strings

Array Members: Maximum number of 100 items.

Length Constraints: Minimum length of 0. Maximum length of 4096.

Pattern: `^[^\u0000\u0085\u2028\u2029\x\n]{0,4096}$`

Required: No

Report

Provides a report detailing the data repository task results of the files processed that match the criteria specified in the report Scope parameter. FSx delivers the report to the file system's linked data repository in Amazon S3, using the path specified in the report Path parameter. You can specify whether or not a report gets generated for a task using the Enabled parameter.

Type: [CompletionReport \(p. 72\)](#) object

Required: No

ResourceARN

The Amazon Resource Name (ARN) for a given resource. ARNs uniquely identify AWS resources. We require an ARN when you need to specify a resource unambiguously across all of AWS. For more information, see [Amazon Resource Names \(ARNs\) and AWS Service Namespaces](#) in the *AWS General Reference*.

Type: String

Length Constraints: Minimum length of 8. Maximum length of 512.

Pattern: `^arn:(?=[^:]+:fsx:[^:]+\d{12}:)(\(|(?=[a-z0-9-.\]{1,63})(?!\\d{1,3}(\.\\d{1,3}){3})(?![^\:]*-{2})(?![^\:]*-\.)(?![^\:]*\.-)[a-z0-9].*(?!-)):\d{4}(?!/)\. {0,1024}$`

Required: No

StartTime

The time that Amazon FSx began processing the task.

Type: Timestamp

Required: No

Status

Provides the status of the number of files that the task has processed successfully and failed to process.

Type: [DataRepositoryTaskStatus](#) (p. 89) object

Required: No

Tags

A list of `Tag` values, with a maximum of 50 elements.

Type: Array of [Tag](#) (p. 109) objects

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

Required: No

TaskId

The system-generated, unique 17-digit ID of the data repository task.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 128.

Pattern: `^(task-[0-9a-f]{17,})$`

Required: Yes

Type

The type of data repository task; `EXPORT_TO_REPOSITORY` is the only type currently supported.

Type: String

Valid Values: `EXPORT_TO_REPOSITORY`

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

DataRepositoryTaskFailureDetails

Provides information about why a data repository task failed. Only populated when the task `Lifecycle` is set to `FAILED`.

Contents

Message

A detailed error message.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

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

DataRepositoryTaskFilter

(Optional) An array of filter objects you can use to filter the response of data repository tasks you will see in the the response. You can filter the tasks returned in the response by one or more file system IDs, task lifecycles, and by task type. A filter object consists of a filter `Name`, and one or more `Values` for the filter.

Contents

Name

Name of the task property to use in filtering the tasks returned in the response.

- Use `file-system-id` to retrieve data repository tasks for specific file systems.
- Use `task-lifecycle` to retrieve data repository tasks with one or more specific lifecycle states, as follows: CANCELED, EXECUTING, FAILED, PENDING, and SUCCEEDED.

Type: String

Valid Values: `file-system-id` | `task-lifecycle`

Required: No

Values

Use Values to include the specific file system IDs and task lifecycle states for the filters you are using.

Type: Array of strings

Array Members: Maximum number of 20 items.

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

Pattern: `^[0-9a-zA-Z*\.\ \\/\?\-_\]*\$`

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

DataRepositoryTaskStatus

Provides the task status showing a running total of the total number of files to be processed, the number successfully processed, and the number of files the task failed to process.

Contents

FailedCount

A running total of the number of files that the task failed to process.

Type: Long

Required: No

LastUpdatedTime

The time at which the task status was last updated.

Type: Timestamp

Required: No

SucceededCount

A running total of the number of files that the task has successfully processed.

Type: Long

Required: No

TotalCount

The total number of files that the task will process. While a task is executing, the sum of `SucceededCount` plus `FailedCount` may not equal `TotalCount`. When the task is complete, `TotalCount` equals the sum of `SucceededCount` plus `FailedCount`.

Type: Long

Required: No

See Also

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

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

DeleteFileSystemLustreConfiguration

The configuration object for the Amazon FSx for Lustre file system being deleted in the `DeleteFileSystem` operation.

Contents

FinalBackupTags

Use if `SkipFinalBackup` is set to `false`, and you want to apply an array of tags to the final backup. If you have set the file system property `CopyTagsToBackups` to `true`, and you specify one or more `FinalBackupTags` when deleting a file system, Amazon FSx will not copy any existing file system tags to the backup.

Type: Array of [Tag \(p. 109\)](#) objects

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

Required: No

SkipFinalBackup

Set `SkipFinalBackup` to `false` if you want to take a final backup of the file system you are deleting. By default, Amazon FSx will not take a final backup on your behalf when the `DeleteFileSystem` operation is invoked. (Default = `true`)

Type: Boolean

Required: No

See Also

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

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

DeleteFileSystemLustreResponse

The response object for the Amazon FSx for Lustre file system being deleted in the `DeleteFileSystem` operation.

Contents

FinalBackupId

The ID of the final backup for this file system.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 128.

Pattern: `^(backup-[0-9a-f]{8,})$`

Required: No

FinalBackupTags

The set of tags applied to the final backup.

Type: Array of [Tag \(p. 109\)](#) objects

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

Required: No

See Also

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

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

DeleteFileSystemWindowsConfiguration

The configuration object for the Microsoft Windows file system used in the `DeleteFileSystem` operation.

Contents

FinalBackupTags

A set of tags for your final backup.

Type: Array of [Tag \(p. 109\)](#) objects

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

Required: No

SkipFinalBackup

By default, Amazon FSx for Windows takes a final backup on your behalf when the `DeleteFileSystem` operation is invoked. Doing this helps protect you from data loss, and we highly recommend taking the final backup. If you want to skip this backup, use this flag to do so.

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

DeleteFileSystemWindowsResponse

The response object for the Microsoft Windows file system used in the `DeleteFileSystem` operation.

Contents

FinalBackupId

The ID of the final backup for this file system.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 128.

Pattern: `^(backup-[0-9a-f]{8,})$`

Required: No

FinalBackupTags

The set of tags applied to the final backup.

Type: Array of [Tag \(p. 109\)](#) objects

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

Required: No

See Also

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

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

FileSystem

A description of a specific Amazon FSx file system.

Contents

AdministrativeActions

A list of administrative actions for the file system that are in process or waiting to be processed. Administrative actions describe changes to the Windows file system that you have initiated using the `UpdateFileSystem` action.

Type: Array of [AdministrativeAction \(p. 65\)](#) objects

Array Members: Maximum number of 50 items.

Required: No

CreationTime

The time that the file system was created, in seconds (since 1970-01-01T00:00:00Z), also known as Unix time.

Type: Timestamp

Required: No

DNSName

The DNS name for the file system.

Type: String

Length Constraints: Minimum length of 16. Maximum length of 275.

Pattern: `^(fsi?-[0-9a-f]{8,})\.{4,253}$`

Required: No

FailureDetails

A structure providing details of any failures that occur when creating the file system has failed.

Type: [FileSystemFailureDetails \(p. 98\)](#) object

Required: No

FileSystemId

The system-generated, unique 17-digit ID of the file system.

Type: String

Length Constraints: Minimum length of 11. Maximum length of 21.

Pattern: `^(fs-[0-9a-f]{8,})$`

Required: No

FileSystemType

The type of Amazon FSx file system, either `LUSTRE` or `WINDOWS`.

Type: String

Valid Values: WINDOWS | LUSTRE

Required: No

KmsKeyId

The ID of the AWS Key Management Service (AWS KMS) key used to encrypt the file system's data for Amazon FSx for Windows File Server file systems and persistent Amazon FSx for Lustre file systems at rest. In either case, if not specified, the Amazon FSx managed key is used. The scratch Amazon FSx for Lustre file systems are always encrypted at rest using Amazon FSx managed keys. For more information, see [Encrypt](#) in the *AWS Key Management Service API Reference*.

Type: String

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

Pattern: `^\{1,2048\}$`

Required: No

Lifecycle

The lifecycle status of the file system, following are the possible values and what they mean:

- **AVAILABLE** - The file system is in a healthy state, and is reachable and available for use.
- **CREATING** - Amazon FSx is creating the new file system.
- **DELETING** - Amazon FSx is deleting an existing file system.
- **FAILED** - An existing file system has experienced an unrecoverable failure. When creating a new file system, Amazon FSx was unable to create the file system.
- **MISCONFIGURED** indicates that the file system is in a failed but recoverable state.
- **UPDATING** indicates that the file system is undergoing a customer initiated update.

Type: String

Valid Values: AVAILABLE | CREATING | FAILED | DELETING | MISCONFIGURED | UPDATING

Required: No

LustreConfiguration

The configuration for the Amazon FSx for Lustre file system.

Type: [LustreFileSystemConfiguration](#) (p. 100) object

Required: No

NetworkInterfaceIds

The IDs of the elastic network interface from which a specific file system is accessible. The elastic network interface is automatically created in the same VPC that the Amazon FSx file system was created in. For more information, see [Elastic Network Interfaces](#) in the *Amazon EC2 User Guide*.

For an Amazon FSx for Windows File Server file system, you can have one network interface ID. For an Amazon FSx for Lustre file system, you can have more than one.

Type: Array of strings

Array Members: Maximum number of 50 items.

Length Constraints: Minimum length of 12. Maximum length of 21.

Pattern: `^(eni-[0-9a-f]{8,})$`

Required: No

OwnerId

The AWS account that created the file system. If the file system was created by an AWS Identity and Access Management (IAM) user, the AWS account to which the IAM user belongs is the owner.

Type: String

Length Constraints: Fixed length of 12.

Pattern: `^\d{12}$`

Required: No

ResourceARN

The Amazon Resource Name (ARN) for the file system resource.

Type: String

Length Constraints: Minimum length of 8. Maximum length of 512.

Pattern: `^arn:(?=[^:]+:fsx:[^:]+\d{12}:)(\(|(?=[a-z0-9-]{1,63})(?!\\d{1,3}\\.\\d{1,3}){3})(?![^\:]*-){2})(?![^\:]*\\.)(?![^\:]*\.-)[a-z0-9].*(?!-))){4}(?!/).\{0,1024}$`

Required: No

StorageCapacity

The storage capacity of the file system in gigabytes (GB).

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 2147483647.

Required: No

StorageType

The storage type of the file system. Valid values are `SSD` and `HDD`. If set to `SSD`, the file system uses solid state drive storage. If set to `HDD`, the file system uses hard disk drive storage.

Type: String

Valid Values: `SSD` | `HDD`

Required: No

SubnetIds

Specifies the IDs of the subnets that the file system is accessible from. For Windows `MULTI_AZ_1` file system deployment type, there are two subnet IDs, one for the preferred file server and one for the standby file server. The preferred file server subnet identified in the `PreferredSubnetID` property. All other file systems have only one subnet ID.

For Lustre file systems, and Single-AZ Windows file systems, this is the ID of the subnet that contains the endpoint for the file system. For `MULTI_AZ_1` Windows file systems, the endpoint for the file system is available in the `PreferredSubnetID`.

Type: Array of strings

Array Members: Maximum number of 50 items.

Length Constraints: Minimum length of 15. Maximum length of 24.

Pattern: `^(subnet-[0-9a-f]{8,})$`

Required: No

Tags

The tags to associate with the file system. For more information, see [Tagging Your Amazon EC2 Resources](#) in the *Amazon EC2 User Guide*.

Type: Array of [Tag \(p. 109\)](#) objects

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

Required: No

VpcId

The ID of the primary VPC for the file system.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 21.

Pattern: `^(vpc-[0-9a-f]{8,})$`

Required: No

WindowsConfiguration

The configuration for this Microsoft Windows file system.

Type: [WindowsFileSystemConfiguration \(p. 114\)](#) object

Required: No

See Also

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

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

FileSystemFailureDetails

A structure providing details of any failures that occur when creating the file system has failed.

Contents

Message

A message describing any failures that occurred during file system creation.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

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

Filter

A filter used to restrict the results of describe calls. You can use multiple filters to return results that meet all applied filter requirements.

Contents

Name

The name for this filter.

Type: String

Valid Values: `file-system-id` | `backup-type` | `file-system-type`

Required: No

Values

The values of the filter. These are all the values for any of the applied filters.

Type: Array of strings

Array Members: Maximum number of 20 items.

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

Pattern: `^[0-9a-zA-Z*\.\ \/ \? \- _]*$`

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

LustreFileSystemConfiguration

The configuration for the Amazon FSx for Lustre file system.

Contents

AutomaticBackupRetentionDays

The number of days to retain automatic backups. Setting this to 0 disables automatic backups. You can retain automatic backups for a maximum of 90 days. The default is 0.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 90.

Required: No

CopyTagsToBackups

A boolean flag indicating whether tags on the file system should be copied to backups. If it's set to true, all tags on the file system are copied to all automatic backups and any user-initiated backups where the user doesn't specify any tags. If this value is true, and you specify one or more tags, only the specified tags are copied to backups. If you specify one or more tags when creating a user-initiated backup, no tags are copied from the file system, regardless of this value. (Default = false)

Type: Boolean

Required: No

DailyAutomaticBackupStartTime

A recurring daily time, in the format `HH:MM`. `HH` is the zero-padded hour of the day (0-23), and `MM` is the zero-padded minute of the hour. For example, `05:00` specifies 5 AM daily.

Type: String

Length Constraints: Fixed length of 5.

Pattern: `^([01]\d|2[0-3]):?([0-5]\d)$`

Required: No

DataRepositoryConfiguration

The data repository configuration object for Lustre file systems returned in the response of the `CreateFileSystem` operation.

Type: [DataRepositoryConfiguration \(p. 81\)](#) object

Required: No

DeploymentType

The deployment type of the FSX for Lustre file system. *Scratch deployment type* is designed for temporary storage and shorter-term processing of data.

`SCRATCH_1` and `SCRATCH_2` deployment types are best suited for when you need temporary storage and shorter-term processing of data. The `SCRATCH_2` deployment type provides in-transit encryption of data and higher burst throughput capacity than `SCRATCH_1`.

The `PERSISTENT_1` deployment type is used for longer-term storage and workloads and encryption of data in transit. To learn more about deployment types, see [FSx for Lustre Deployment Options](#). (Default = `SCRATCH_1`)

Type: String

Valid Values: `SCRATCH_1` | `SCRATCH_2` | `PERSISTENT_1`

Required: No

DriveCacheType

The type of drive cache used by `PERSISTENT_1` file systems that are provisioned with HDD storage devices. This parameter is required when storage type is HDD. Set to `READ`, improve the performance for frequently accessed files and allows 20% of the total storage capacity of the file system to be cached.

This parameter is required when `StorageType` is set to `HDD`.

Type: String

Valid Values: `NONE` | `READ`

Required: No

MountName

You use the `MountName` value when mounting the file system.

For the `SCRATCH_1` deployment type, this value is always `fsx`. For `SCRATCH_2` and `PERSISTENT_1` deployment types, this value is a string that is unique within an AWS Region.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 8.

Pattern: `^[A-Za-z0-9_-]{1,8}$`

Required: No

PerUnitStorageThroughput

Per unit storage throughput represents the megabytes per second of read or write throughput per 1 tebibyte of storage provisioned. File system throughput capacity is equal to `Storage capacity (TiB) * PerUnitStorageThroughput (MB/s/TiB)`. This option is only valid for `PERSISTENT_1` deployment types.

Valid values for SSD storage: 50, 100, 200. Valid values for HDD storage: 12, 40.

Type: Integer

Valid Range: Minimum value of 12. Maximum value of 200.

Required: No

WeeklyMaintenanceStartTime

The preferred start time to perform weekly maintenance, formatted `d:HH:MM` in the UTC time zone. `d` is the weekday number, from 1 through 7, beginning with Monday and ending with Sunday.

Type: String

Length Constraints: Fixed length of 7.

Pattern: `^[1-7]:([01]\d|2[0-3]):?([0-5]\d)$`

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

SelfManagedActiveDirectoryAttributes

The configuration of the self-managed Microsoft Active Directory (AD) directory to which the Windows File Server instance is joined.

Contents

DnsIps

A list of up to two IP addresses of DNS servers or domain controllers in the self-managed AD directory.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 2 items.

Length Constraints: Minimum length of 7. Maximum length of 15.

Pattern: `^(([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.){3}([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])$`

Required: No

DomainName

The fully qualified domain name of the self-managed AD directory.

Type: String

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

Pattern: `^[^\u0000\u0085\u2028\u2029\x\n]{1,255}$`

Required: No

FileSystemAdministratorsGroup

The name of the domain group whose members have administrative privileges for the FSx file system.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^[^\u0000\u0085\u2028\u2029\x\n]{1,256}$`

Required: No

OrganizationalUnitDistinguishedName

The fully qualified distinguished name of the organizational unit within the self-managed AD directory to which the Windows File Server instance is joined.

Type: String

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

Pattern: `^[^\u0000\u0085\u2028\u2029\x\n]{1,2000}$`

Required: No

UserName

The user name for the service account on your self-managed AD domain that FSx uses to join to your AD domain.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^[^\u0000\u0085\u2028\u2029\r\n]{1,256}$`

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

SelfManagedActiveDirectoryConfiguration

The configuration that Amazon FSx uses to join the Windows File Server instance to your self-managed (including on-premises) Microsoft Active Directory (AD) directory.

Contents

DnsIps

A list of up to two IP addresses of DNS servers or domain controllers in the self-managed AD directory. The IP addresses need to be either in the same VPC CIDR range as the one in which your Amazon FSx file system is being created, or in the private IP version 4 (IPv4) address ranges, as specified in [RFC 1918](#):

- 10.0.0.0 - 10.255.255.255 (10/8 prefix)
- 172.16.0.0 - 172.31.255.255 (172.16/12 prefix)
- 192.168.0.0 - 192.168.255.255 (192.168/16 prefix)

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 2 items.

Length Constraints: Minimum length of 7. Maximum length of 15.

Pattern: `^((([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.){3}([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5]))$`

Required: Yes

DomainName

The fully qualified domain name of the self-managed AD directory, such as `corp.example.com`.

Type: String

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

Pattern: `^[^\\u0000\u0085\u2028\u2029\x\n]{1,255}$`

Required: Yes

FileSystemAdministratorsGroup

(Optional) The name of the domain group whose members are granted administrative privileges for the file system. Administrative privileges include taking ownership of files and folders, setting audit controls (audit ACLs) on files and folders, and administering the file system remotely by using the FSx Remote PowerShell. The group that you specify must already exist in your domain. If you don't provide one, your AD domain's Domain Admins group is used.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^[^\\u0000\u0085\u2028\u2029\x\n]{1,256}$`

Required: No

OrganizationalUnitDistinguishedName

(Optional) The fully qualified distinguished name of the organizational unit within your self-managed AD directory that the Windows File Server instance will join. Amazon FSx only accepts OU

as the direct parent of the file system. An example is `OU=FSx,DC=yourdomain,DC=corp,DC=com`. To learn more, see [RFC 2253](#). If none is provided, the FSx file system is created in the default location of your self-managed AD directory.

Important

Only Organizational Unit (OU) objects can be the direct parent of the file system that you're creating.

Type: String

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

Pattern: `^[^\u0000\u0085\u2028\u2029\r\n]{1,2000}$`

Required: No

Password

The password for the service account on your self-managed AD domain that Amazon FSx will use to join to your AD domain.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^.{1,256}$`

Required: Yes

UserName

The user name for the service account on your self-managed AD domain that Amazon FSx will use to join to your AD domain. This account must have the permission to join computers to the domain in the organizational unit provided in `OrganizationalUnitDistinguishedName`, or in the default location of your AD domain.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^[^\u0000\u0085\u2028\u2029\r\n]{1,256}$`

Required: Yes

See Also

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

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

SelfManagedActiveDirectoryConfigurationUpdates

The configuration that Amazon FSx uses to join the Windows File Server instance to a self-managed Microsoft Active Directory (AD) directory.

Contents

DnsIps

A list of up to two IP addresses of DNS servers or domain controllers in the self-managed AD directory.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 2 items.

Length Constraints: Minimum length of 7. Maximum length of 15.

Pattern: `^((([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.){3}([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5]))$`

Required: No

Password

The password for the service account on your self-managed AD domain that Amazon FSx will use to join to your AD domain.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^\.{1,256}$`

Required: No

UserName

The user name for the service account on your self-managed AD domain that Amazon FSx will use to join to your AD domain. This account must have the permission to join computers to the domain in the organizational unit provided in `OrganizationalUnitDistinguishedName`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^[^\u0000\u0085\u2028\u2029\r\n]{1,256}$`

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

Tag

Specifies a key-value pair for a resource tag.

Contents

Key

A value that specifies the `TagKey`, the name of the tag. Tag keys must be unique for the resource to which they are attached.

Type: String

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

Pattern: `^[a-zA-Z0-9_./+=\-\@]*$`

Required: Yes

Value

A value that specifies the `TagValue`, the value assigned to the corresponding tag key. Tag values can be null and don't have to be unique in a tag set. For example, you can have a key-value pair in a tag set of `finances : April` and also of `payroll : April`.

Type: String

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

Pattern: `^[a-zA-Z0-9_./+=\-\@]*$`

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

UpdateFileSystemLustreConfiguration

The configuration object for Amazon FSx for Lustre file systems used in the `UpdateFileSystem` operation.

Contents

AutoImportPolicy

(Optional) When you create your file system, your existing S3 objects appear as file and directory listings. Use this property to choose how Amazon FSx keeps your file and directory listing up to date as you add or modify objects in your linked S3 bucket. `AutoImportPolicy` can have the following values:

- `NONE` - (Default) `AutoImport` is off. Amazon FSx only updates file and directory listings from the linked S3 bucket when the file system is created. FSx does not update the file and directory listing for any new or changed objects after choosing this option.
- `NEW` - `AutoImport` is on. Amazon FSx automatically imports directory listings of any new objects added to the linked S3 bucket that do not currently exist in the FSx file system.
- `NEW_CHANGED` - `AutoImport` is on. Amazon FSx automatically imports file and directory listings of any new objects added to the S3 bucket and any existing objects that are changed in the S3 bucket after you choose this option.

For more information, see [Automatically import updates from your S3 bucket](#).

Type: String

Valid Values: `NONE` | `NEW` | `NEW_CHANGED`

Required: No

AutomaticBackupRetentionDays

The number of days to retain automatic backups. Setting this to 0 disables automatic backups. You can retain automatic backups for a maximum of 90 days. The default is 0.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 90.

Required: No

DailyAutomaticBackupStartTime

A recurring daily time, in the format `HH:MM`. `HH` is the zero-padded hour of the day (0-23), and `MM` is the zero-padded minute of the hour. For example, `05:00` specifies 5 AM daily.

Type: String

Length Constraints: Fixed length of 5.

Pattern: `^([01]\d|2[0-3]):?([0-5]\d)$`

Required: No

WeeklyMaintenanceStartTime

(Optional) The preferred start time to perform weekly maintenance, formatted `d:HH:MM` in the UTC time zone. `d` is the weekday number, from 1 through 7, beginning with Monday and ending with Sunday.

Type: String

Length Constraints: Fixed length of 7.

Pattern: `^[1-7]:([01]\d|2[0-3]):?([0-5]\d)$`

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

UpdateFileSystemWindowsConfiguration

Updates the configuration for an existing Amazon FSx for Windows File Server file system. Amazon FSx only overwrites existing properties with non-null values provided in the request.

Contents

AutomaticBackupRetentionDays

The number of days to retain automatic daily backups. Setting this to zero (0) disables automatic daily backups. You can retain automatic daily backups for a maximum of 90 days. For more information, see [Working with Automatic Daily Backups](#).

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 90.

Required: No

DailyAutomaticBackupStartTime

The preferred time to start the daily automatic backup, in the UTC time zone, for example, 02:00

Type: String

Length Constraints: Fixed length of 5.

Pattern: `^([01]\d|2[0-3]):?([0-5]\d)$`

Required: No

SelfManagedActiveDirectoryConfiguration

The configuration Amazon FSx uses to join the Windows File Server instance to the self-managed Microsoft AD directory. You cannot make a self-managed Microsoft AD update request if there is an existing self-managed Microsoft AD update request in progress.

Type: [SelfManagedActiveDirectoryConfigurationUpdates \(p. 107\)](#) object

Required: No

ThroughputCapacity

Sets the target value for a file system's throughput capacity, in MB/s, that you are updating the file system to. Valid values are 8, 16, 32, 64, 128, 256, 512, 1024, 2048. You cannot make a throughput capacity update request if there is an existing throughput capacity update request in progress. For more information, see [Managing Throughput Capacity](#).

Type: Integer

Valid Range: Minimum value of 8. Maximum value of 2048.

Required: No

WeeklyMaintenanceStartTime

The preferred start time to perform weekly maintenance, formatted d:HH:MM in the UTC time zone. Where d is the weekday number, from 1 through 7, with 1 = Monday and 7 = Sunday.

Type: String

Length Constraints: Fixed length of 7.

Pattern: `^[1-7]:([01]\d|2[0-3]):?([0-5]\d)$`

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

WindowsFileSystemConfiguration

The configuration for this Microsoft Windows file system.

Contents

ActiveDirectoryId

The ID for an existing Microsoft Active Directory instance that the file system should join when it's created.

Type: String

Length Constraints: Fixed length of 12.

Pattern: `^d-[0-9a-f]{10}$`

Required: No

AutomaticBackupRetentionDays

The number of days to retain automatic backups. Setting this to 0 disables automatic backups. You can retain automatic backups for a maximum of 90 days.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 90.

Required: No

CopyTagsToBackups

A boolean flag indicating whether tags on the file system should be copied to backups. This value defaults to false. If it's set to true, all tags on the file system are copied to all automatic backups and any user-initiated backups where the user doesn't specify any tags. If this value is true, and you specify one or more tags, only the specified tags are copied to backups. If you specify one or more tags when creating a user-initiated backup, no tags are copied from the file system, regardless of this value.

Type: Boolean

Required: No

DailyAutomaticBackupStartTime

The preferred time to take daily automatic backups, in the UTC time zone.

Type: String

Length Constraints: Fixed length of 5.

Pattern: `^([01]\d|2[0-3]):?([0-5]\d)$`

Required: No

DeploymentType

Specifies the file system deployment type, valid values are the following:

- `MULTI_AZ_1` - Specifies a high availability file system that is configured for Multi-AZ redundancy to tolerate temporary Availability Zone (AZ) unavailability, and supports SSD and HDD storage.
- `SINGLE_AZ_1` - (Default) Specifies a file system that is configured for single AZ redundancy, only supports SSD storage.

- `SINGLE_AZ_2` - Latest generation Single AZ file system. Specifies a file system that is configured for single AZ redundancy and supports SSD and HDD storage.

For more information, see [Single-AZ and Multi-AZ File Systems](#).

Type: String

Valid Values: `MULTI_AZ_1` | `SINGLE_AZ_1` | `SINGLE_AZ_2`

Required: No

MaintenanceOperationsInProgress

The list of maintenance operations in progress for this file system.

Type: Array of strings

Array Members: Maximum number of 20 items.

Valid Values: `PATCHING` | `BACKING_UP`

Required: No

PreferredFileServerIp

For `MULTI_AZ_1` deployment types, the IP address of the primary, or preferred, file server.

Use this IP address when mounting the file system on Linux SMB clients or Windows SMB clients that are not joined to a Microsoft Active Directory. Applicable for all Windows file system deployment types. This IP address is temporarily unavailable when the file system is undergoing maintenance. For Linux and Windows SMB clients that are joined to an Active Directory, use the file system's DNSName instead. For more information on mapping and mounting file shares, see [Accessing File Shares](#).

Type: String

Length Constraints: Minimum length of 7. Maximum length of 15.

Pattern: `^(([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.){3}([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])$`

Required: No

PreferredSubnetId

For `MULTI_AZ_1` deployment types, it specifies the ID of the subnet where the preferred file server is located. Must be one of the two subnet IDs specified in `SubnetIds` property. Amazon FSx serves traffic from this subnet except in the event of a failover to the secondary file server.

For `SINGLE_AZ_1` and `SINGLE_AZ_2` deployment types, this value is the same as that for `SubnetIDs`. For more information, see [Availability and Durability: Single-AZ and Multi-AZ File Systems](#)

Type: String

Length Constraints: Minimum length of 15. Maximum length of 24.

Pattern: `^(subnet-[0-9a-f]{8,})$`

Required: No

RemoteAdministrationEndpoint

For `MULTI_AZ_1` deployment types, use this endpoint when performing administrative tasks on the file system using Amazon FSx Remote PowerShell.

For `SINGLE_AZ_1` and `SINGLE_AZ_2` deployment types, this is the DNS name of the file system.

This endpoint is temporarily unavailable when the file system is undergoing maintenance.

Type: String

Length Constraints: Minimum length of 16. Maximum length of 275.

Pattern: `^(fsi?-[0-9a-f]{8,})\.{4,253}$`

Required: No

SelfManagedActiveDirectoryConfiguration

The configuration of the self-managed Microsoft Active Directory (AD) directory to which the Windows File Server instance is joined.

Type: [SelfManagedActiveDirectoryAttributes \(p. 103\)](#) object

Required: No

ThroughputCapacity

The throughput of an Amazon FSx file system, measured in megabytes per second.

Type: Integer

Valid Range: Minimum value of 8. Maximum value of 2048.

Required: No

WeeklyMaintenanceStartTime

The preferred start time to perform weekly maintenance, formatted `d:HH:MM` in the UTC time zone. `d` is the weekday number, from 1 through 7, beginning with Monday and ending with Sunday.

Type: String

Length Constraints: Fixed length of 7.

Pattern: `^[1-7]:([01]\d|2[0-3]):?([0-5]\d)$`

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 Go](#)
- [AWS SDK for Java](#)
- [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 [Signature Version 4 Signing Process](#) in the *Amazon Web Services General Reference*.

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 [Task 2: Create a String to Sign for Signature Version 4](#) in the *Amazon Web Services General Reference*.

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'THHMMSS'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 [Handling Dates in Signature Version 4](#) in the *Amazon Web Services General Reference*.

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 Security Token Service, go to [AWS Services That Work with IAM](#) in the *IAM User Guide*.

Condition: If you're using temporary security credentials from the AWS Security Token Service, 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 [Task 1: Create a Canonical Request For Signature Version 4](#) in the *Amazon Web Services General Reference*.

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