## Amazon Elastic File System User Guide



#### Amazon Elastic File System: User Guide

Copyright © 2016 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

What Is Amazon Elastic File System?	. 1
Are you a first-time user of Amazon EFS?	. 1
How it Works	. 3
Overview	. 3
How Amazon EFS Works with Amazon EC2	. 4
How Amazon EFS Works with AWS Direct Connect	. 4
Implementation Summary	5
Authentication and Access Control	6
Setting I In	. 0 8
Sign up for AWS	. 0 8
Create an IAM Lear	. 0 8
Cetting Started	10
Accumptione	10
Assumptions	10
Related Topics	11
Step 1: Create Your EC2 Resources and Launch Your EC2 Instance	11
Step 2: Create Your Amazon EFS File System	15
Step 3: Connect to Your Amazon EC2 Instance and Mount the Amazon EFS File System	16
Step 4: Clean Up Resources and Protect Your AWS Account	17
Creating Resources for Amazon EFS	18
Creating File Systems	19
Requirements	19
Permissions Required	20
Creating a File System	20
Creating Mount Targets	23
Creating a Mount Target Using the Amazon EFS console	24
Creating a Mount Target using the AWS CLI	27
Creating Security Groups	28
Creating Security Groups Using the AWS Management Console	28
Creating Security Groups Using the AWS CLI	29
Managing File Systems	30
Managing Network Accessibility	30
Creating or Deleting Mount Tarrets in a VPC	32
Creating of Detering Would rargets in a Vio	24
Undeting the Mount Target Configuration	25
Appending the Mount Target Configuration	30
Managing Tags	31
	37
Using the AWS CLI	37
Deleting a File System	37
Using the Console	37
Using the CLI	38
Related Topics	38
Mounting File Systems	39
NFS Support	39
Troubleshooting AMI/Kernel Versions	39
Installing the NFS Client	40
Mounting on Amazon EC2 with a DNS Name	41
Mounting on On-Premises Servers with a DNS Name	42
Mounting with an IP Address	42
Mounting Automatically	43
Updating an Existing EC2 Instance to Mount Automatically	43
Configuring an EES File System to Mount Automatically at EC2 Instance Launch	45
Additional Mounting Considerations	46
Unmounting File Systems	46
Monitoring File Systems	48
Monitoring Tools	<u>⊿</u> 8
	10

Automated Tools	. 49
Manual Monitoring Tools	49
Monitoring CloudWatch	49
Amazon CloudWatch Metrics for Amazon EFS	50
Amazon EFS Dimensions	. 52
How Do I Use Amazon EFS Metrics?	. 52
Access CloudWatch Metrics	. 53
Creating Alarms	53
Using File Systems	55
Data Consistency in Amazon EFS	55
Related Topics	. 55
NFS-Level Users. Groups. and Permissions	56
Example Amazon EFS File System Use Cases and Permissions	56
User and group ID permissions on files and directories within a file system	57
No Root Squashing	58
Permissions Caching	58
Changing Ello System Object Ownership	50
	50
Metering File System and Object Sizes	58
Metering Amazon EFS File System Objects	58
Metering an Amazon EFS File System	59
Unsupported NFSv4 Features	60
Performance	. 62
Performance Overview	62
Amazon EFS Use Cases	. 63
Big Data and Analytics	. 63
Media Processing Workflows	63
Content Management and Web Serving	. 63
Home Directories	63
Performance Modes	63
General Purpose Performance Mode	. 64
Max I/O Performance Mode	64
Using the Right Performance Mode	64
Bursting	64
Managing Burst Credite	66
On Promises Deformance Considerations	66
Architectus for Lish Aveilebility	67
	. 07
Anazon Ers renormance rips	07
Related Topics	. 68
	69
AWS Identity and Access Management (IAM) permissions for API calls	69
Security Groups for EC2 Instances and Mount Targets	69
Security considerations for mounting an Amazon EFS file system	70
Read, Write, and Execute Permissions for Files and Directories	71
Limits	72
Amazon EFS Limits That Can Be Increased	. 72
Amazon EFS Resource Limits	. 72
Limits for Client EC2 Instances	73
Limits for Amazon EFS File Systems	73
Additional Considerations	73
Troubleshooting Amazon EFS	. 75
Troubleshooting General Issues	75
Mount Command Fails with "wrong is type" Error Message	76
Mount Command Fails with "incorrect mount option" Error Message	76
File System Mount Hange and Then Faile with Timesut Free	70
File System Mount Heing Mount Terret DNC Name Faile	10 77
	11
EU2 Instance Hangs	. //
Mount Larget Litecycle State Is Stuck	77
File System Mount on Windows Instance Fails	77

Application Writing Large Amounts of Data Hangs	. 78
Mount Does Not Respond	78
Open and Close Operations Are Serialized	78
Operations on Newly Mounted File System Return "bad file handle" Error	. 79
Custom NFS Settings Causing Write Delays	79
File Operation Errors	80
Command Fails with "Disk guota exceeded" Error	80
Command Fails with "I/O error"	80
Command Fails with "File name is too long" Error	80
Command Fails with "Too many links" Error	80
Command Fails with "File too large" Error	Q1
	01
Command Fails with Try again End	. 01
Houbleshouling Almi and Kemel Names	. 01
Unable to chown on Kernel Version 2.6.32	. 81
File System Keeps Performing Operations Repeatedly Due to Client Bug	81
Deadlocked Client	82
Walkthroughs	. 83
Walkthrough 1: Create and Mount a File System Using the AWS CLI	83
Before You Begin	. 84
Setting Up Tools	. 84
Step 1: Create Amazon EC2 Resources	85
Step 2: Create Amazon EFS Resources	. 89
Step 3: Mount and Test the File System	. 92
Step 4: Clean up	94
Walkthrough 2: Set Up an Apache Web Server and Serve Files	96
Single EC2 Instance Serving Files	. 96
Multiple EC2 Instances Serving Files	98
Walkthrough 3: Create Writable Per-User Subdirectories	101
Automatic Remounting on Reboot	102
Walkthrough 4 <sup>·</sup> Back Up an FES File System	102
FES Backup Performance	103
Considerations for EES Backup	103
Assumptions for EES Backup	103
Recking up an Amazon EES File System	104
Additional Decourses	104
Multiplier Constant and Maunt a File System On Promised with AW/S Direct Connect	109
Potoro Vou Bogin	114
Stop 1: Croote Vour Amazon Electic File System Resources	114
Step 1. Create Four Annazon EES File System as Vour On Dramines Server	110
Step 2: Nount the Anazon Ers File System on Four On-Fremises Server	110
Step 3. Clean Op Resources and Protect Your Avvs Account	117
Authentication and Access Control	118
Autnentication	118
Access Control	119
Overview of Managing Access	120
Amazon Elastic File System Resources and Operations	120
Understanding Resource Ownership	120
Managing Access to Resources	121
Specifying Policy Elements: Actions, Effects, and Principals	122
Specifying Conditions in a Policy	123
Using Identity-Based Policies (IAM Policies)	123
Permissions Required to Use the Amazon Elastic File System Console	124
AWS Managed (Predefined) Policies for Amazon Elastic File System	125
Customer Managed Policy Examples	125
Amazon EFS API Permissions Reference	127
Amazon EFS API	129
API Endpoint	129
API Version	130
Related Topics	130

Actions 1	30
CreateFileSystem1	31
CreateMountTarget 1	35
CreateTags 1	41
DeleteFileSystem	43
DeleteMountTarget	45
DeleteTags	47
DescribeFileSystems	49
DescribeMountTargets1	52
DescribeMountTargetSecurityGroups1	55
DescribeTags	57
ModifyMountTargetSecurityGroups1	60
Data Types 1	61
FileSystemDescription	63
FileSystemSize	65
MountTargetDescription1	66
Tag 1	67
Document History 1	68

# What Is Amazon Elastic File System?

Amazon Elastic File System (Amazon EFS) provides simple, scalable file storage for use with Amazon EC2. With Amazon EFS, storage capacity is elastic, growing and shrinking automatically as you add and remove files, so your applications have the storage they need, when they need it.

Amazon EFS has a simple web services interface that allows you to create and configure file systems quickly and easily. The service manages all the file storage infrastructure for you, avoiding the complexity of deploying, patching, and maintaining complex file system deployments.

Amazon EFS supports the Network File System version 4.1 (NFSv4.1) protocol, so the applications and tools that you use today work seamlessly with Amazon EFS. Multiple Amazon EC2 instances can access an Amazon EFS file system at the same time, providing a common data source for workloads and applications running on more than one instance or server.

With Amazon EFS, you pay only for the storage used by your file system. You don't need to provision storage in advance and there is no minimum fee or setup cost. For more information, see Amazon EFS Pricing.

The service is designed to be highly scalable, highly available, and highly durable. Amazon EFS file systems store data and metadata across multiple Availability Zones in a region and can grow to petabyte scale, drive high levels of throughput, and allow massively parallel access from Amazon EC2 instances to your data.

Amazon EFS provides file system access semantics, such as strong data consistency and file locking. For more information, see Data Consistency in Amazon EFS (p. 55). Amazon EFS also allows you to tightly control access to your file systems through POSIX permissions.

Amazon EFS is designed to provide the throughput, IOPS, and low latency needed for a broad range of workloads. With Amazon EFS, throughput and IOPS scale as a file system grows, and file operations are delivered with consistent, low latencies. For more information, see Amazon EFS Performance (p. 62).

#### Note

Using Amazon EFS with Microsoft Windows Amazon EC2 instances is not supported.

## Are you a first-time user of Amazon EFS?

If you are a first-time user of Amazon EFS, we recommend you read the following sections in order:

- 1. For an Amazon EFS product and pricing overview, see Amazon EFS.
- 2. For an Amazon EFS technical overview, see Amazon EFS: How It Works (p. 3).
- 3. Try the introductory exercises:
  - Getting Started (p. 10)
  - Walkthroughs (p. 83)

If you would like to learn more about Amazon EFS, the following topics discuss the service in greater detail:

- Creating Resources for Amazon EFS (p. 18)
- Managing Amazon EFS File Systems (p. 30)
- Amazon EFS API (p. 129)

## Amazon EFS: How It Works

Following, you can find a description about how Amazon EFS works, its implementation details, and security considerations.

#### Topics

- Overview (p. 3)
- Implementation Summary (p. 5)

## **Overview**

Amazon EFS provides file storage in the AWS Cloud. With Amazon EFS, you can create a file system, mount the file system on an Amazon EC2 instance, and then read and write data from to and from your file system. You can mount an Amazon EFS file system in your VPC, through the Network File System version 4.1 (NFSv4.1) protocol.

For a list of Amazon EC2 Linux Amazon Machine Images (AMIs) that support this protocol, see NFS Support (p. 39). We recommend using a current generation Linux NFSv4.1 client, such as those found in Amazon Linux and Ubuntu AMIs. For some AMIs, you'll need to install an NFS client to mount your file system on your Amazon EC2 instance. For instructions, see Installing the NFS Client (p. 40).

You can access your Amazon EFS file system concurrently from Amazon EC2 instances in your Amazon VPC, so applications that scale beyond a single connection can access a file system. Amazon EC2 instances running in multiple Availability Zones within the same region can access the file system, so that many users can access and share a common data source.

Note the following restrictions:

- You can mount an Amazon EFS file system on instances in only one VPC at a time.
- Both the file system and VPC must be in the same AWS Region.

For a list of AWS regions where you can create an Amazon EFS file system, see the Amazon Web Services General Reference.

To access your Amazon EFS file system in a VPC, you create one or more *mount targets* in the VPC. You mount your file system from these mount targets. A mount target represents an IP address for an NFSv4.1 endpoint at which you can mount an Amazon EFS file system. You can create one mount target in each of the Availability Zones. If there are multiple subnets in an Availability Zone in your VPC, you create a mount target in one of the subnets, and all EC2 instances in that Availability Zone share that mount target. Mount targets themselves are designed to be highly available. When designing your application for high availability and the ability to failover to other Availability Zones, keep in mind that the IP addresses and DNS for your mount targets in each Availability Zone are static.

After mounting the file system via the mount target, you use it like any other POSIX-compliant file system. For information about NFS-level permissions and related considerations, see Network File System (NFS)–Level Users, Groups, and Permissions (p. 56).

You can mount your Amazon EFS file systems on your on-premises datacenter servers when connected to your Amazon VPC with AWS Direct Connect. You can mount your EFS file systems on on-premises servers to migrate data sets to EFS, enable cloud bursting scenarios, or backup your on-premises data to EFS.

Amazon EFS file systems can be mounted on Amazon EC2 instances, or on-premises through an AWS Direct Connect connection.

### How Amazon EFS Works with Amazon EC2

The following illustration shows an example VPC accessing an Amazon EFS file system. Here, EC2 instances in the VPC have file systems mounted.



In this illustration, the VPC has three Availability Zones, and each has one mount target created in it. We recommend that you access the file system from a mount target within the same Availability Zone. Note that one of the Availability Zones has two subnets. However, a mount target is created in only one of the subnets. Creating this setup works as follows:

- 1. Create your Amazon EC2 resources and launch your Amazon EC2 instance. For more information on Amazon EC2, see Amazon EC2 Virtual Server Hosting.
- 2. Create your Amazon EFS file system.
- 3. Connect to your Amazon EC2 instance, and mount the Amazon EFS file system.

For detailed steps, see Getting Started with Amazon Elastic File System (p. 10).

### How Amazon EFS Works with AWS Direct Connect

By using an Amazon EFS file system mounted on an on-premises server, you can migrate on-premises data into the AWS Cloud hosted in an Amazon EFS file system. You can also take advantage of bursting, meaning that you can move data from your on-premises servers into Amazon EFS, analyze

it on a fleet of Amazon EC2 instances in your Amazon VPC, and then store the results permanently in your file system or move the results back to your on-premises server.

Keep the following considerations in mind when using Amazon EFS with AWS Direct Connect:

- Your on-premises server must have a Linux based operating system. We recommend Linux kernel version 4.0 or later.
- For the sake of simplicity, we recommend mounting an Amazon EFS file system on an on-premises server using a mount target IP address instead of a DNS name.
- AWS VPN is not supported for accessing an Amazon EFS file system from an on-premises server.

There is no additional cost for on-premises access to your Amazon EFS file systems. Note that you'll be charged for the AWS Direct Connect connection to your Amazon VPC. For more information, see AWS Direct Connect Pricing.

The following illustration shows an example of how to access an Amazon EFS file system from onpremises (the on-premises servers have the file systems mounted).



You can use any one of the mount targets in your VPC as long as the subnet of the mount target is reachable by using the AWS Direct Connect connection between your on-premises server and your Amazon VPC. To access Amazon EFS from a on-premises server, you need to add a rule to your mount target security group to allow inbound traffic to the NFS port (2049) from your on-premises server.

To create a setup like this, you do the following:

- 1. Establish an AWS Direct Connect connection between your on-premises data center and your Amazon VPC. For more information on AWS Direct Connect, see AWS Direct Connect.
- 2. Create your Amazon EFS file system.
- 3. Mount the Amazon EFS file system on your on-premises server.

For detailed steps, see Walkthrough 5: Create and Mount a File System On-Premises with AWS Direct Connect (p. 114).

## Implementation Summary

In Amazon EFS, a file system is the primary resource. Each file system has properties such as ID, creation token, creation time, file system size in bytes, number of mount targets created for the file system, and the file system state. For more information, see CreateFileSystem (p. 131).

Amazon EFS also supports other resources to configure the primary resource. These include mount targets and tags:

• Mount target – To access your file system, you must create mount targets in your VPC. Each mount target has the following properties: the mount target ID, the subnet ID in which it is created, the file system ID for which it is created, an IP address at which the file system may be mounted, and the mount target state. You can use the IP address or the DNS name in your mount command. Each mount target has a DNS name of the following form.

availability-zone.file-system-id.efs.aws-region.amazonaws.com

You can specify this DNS name in your mount command to mount the Amazon EFS file system. Suppose you create an ~/efs-mount-point subdirectory on your EC2 instance or on-premises server. Then, you can use the mount command to mount the file system. For example, on an Amazon Linux AMI, you can use following mount command.

```
$ sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 mount-
target-DNS:/ ~/efs-mount-point
```

For more information, see Creating Mount Targets (p. 23). First, you need to install the NFS client on your EC2 instance. The Getting Started (p. 10) exercise provides step-by-step instructions.

• **Tags** – To help organize your file systems, you can assign your own metadata to each of the file systems you create. Each tag is a key-value pair.

You can think of mount targets and tags as *subresources* that don't exist without being associated with a file system.

Amazon EFS provides API operations for you to create and manage these resources. In addition to the create and delete operations for each resource, Amazon EFS also supports a describe operation that enables you to retrieve resource information. You have the following options for creating and managing these resources:

- Use the Amazon EFS console For an example, see Getting Started (p. 10).
- Use the Amazon EFS command line interface (CLI) For an example, see Walkthrough 1: Create Amazon EFS File System and Mount It on an EC2 Instance Using the AWS CLI (p. 83).
- You can also manage these resources programmatically as follows:
  - Use the AWS SDKs The AWS SDKs simplify your programming tasks by wrapping the underlying Amazon EFS API. The SDK clients also authenticate your requests by using access keys that you provide. For more information, see Sample Code and Libraries.
  - Call the Amazon EFS API directly from your application If you cannot use the SDKs for some reason, you can make the Amazon EFS API calls directly from your application. However, you need to write the necessary code to authenticate your requests if you use this option. For more information about the Amazon EFS API, see Amazon EFS API (p. 129).

## Authentication and Access Control

You must have valid credentials to make Amazon EFS API requests, such as create a file system. In addition, you must also have permissions to create or access resources. By default, when you use the root account credentials of your AWS account you can create and access resources owned by that account. However, we do not recommend using root account credentials. In addition, any AWS Identity and Access Management (IAM) users and roles you create in your account must be granted permissions to create or access resources. For more information about permissions, see Authentication and Access Control for Amazon EFS (p. 118).

## Setting Up

Before you use Amazon EFS for the first time, complete the following tasks:

- 1. Sign up for AWS (p. 8)
- 2. Create an IAM User (p. 8)

## Sign up for AWS

When you sign up for Amazon Web Services (AWS), your AWS account is automatically signed up for all services in AWS, including Amazon EFS. You are charged only for the services that you use.

With Amazon EFS, you pay only for the storage you use. For more information about Amazon EFS usage rates, see the Amazon Elastic File System Pricing. If you are a new AWS customer, you can get started with Amazon EFS for free. For more information, see AWS Free Usage Tier.

If you have an AWS account already, skip to the next task. If you don't have an AWS account, use the following procedure to create one.

#### To create an AWS account

- 1. Open http://aws.amazon.com/, and then choose Create an AWS Account.
- 2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call and entering a PIN using the phone keypad.

Note your AWS account number, because you'll need it for the next task.

## Create an IAM User

Services in AWS, such as Amazon EFS, require that you provide credentials when you access them, so that the service can determine whether you have permissions to access its resources. AWS recommends that you do not use the root credentials of your AWS account to make requests. Instead, create an IAM user, and grant that user full access. We refer to these users as administrator users. You can use the administrator user credentials, instead of root credentials of your account, to interact with AWS and perform tasks, such as create a bucket, create users, and grant them permissions.

For more information, see Root Account Credentials vs. IAM User Credentials in the AWS General Reference and IAM Best Practices in the IAM User Guide.

If you signed up for AWS but have not created an IAM user for yourself, you can create one using the IAM console.

#### To create an IAM user for yourself and add the user to an Administrators group

- 1. Sign in to the Identity and Access Management (IAM) console at https://console.aws.amazon.com/ iam/.
- 2. In the navigation pane, choose **Users**, and then choose **Add user**.
- 3. For **User name**, type a user name, such as **Administrator**. The name can consist of letters, digits, and the following characters: plus (+), equal (=), comma (,), period (.), at (@), underscore (\_), and hyphen (-). The name is not case sensitive and can be a maximum of 64 characters in length.
- 4. Select the check box next to **AWS Management Console access**, select **Custom password**, and then type the new user's password in the text box. You can optionally select **Require password reset** to force the user to select a new password the next time the user signs in.
- 5. Choose Next: Permissions.
- 6. On the **Set permissions for user** page, choose **Add user to group**.
- 7. Choose Create group.
- 8. In the **Create group** dialog box, type the name for the new group. The name can consist of letters, digits, and the following characters: plus (+), equal (=), comma (,), period (.), at (@), underscore (\_), and hyphen (-). The name is not case sensitive and can be a maximum of 128 characters in length.
- 9. For Filter, choose Job function.
- 10. In the policy list, select the check box for AdministratorAccess. Then choose Create group.
- 11. Back in the list of groups, select the check box for your new group. Choose **Refresh** if necessary to see the group in the list.
- 12. Choose **Next: Review** to see the list of group memberships to be added to the new user. When you are ready to proceed, choose Add permissions.

You can use this same process to create more groups and users, and to give your users access to your AWS account resources. To learn about using policies to restrict users' permissions to specific AWS resources, go to Access Management and Example Policies for Administering AWS Resources.

To sign in as this new IAM user, sign out of the AWS Management Console, and then use the following URL, where *your\_aws\_account\_id* is your AWS account number without the hyphens (for example, if your AWS account number is 1234-5678-9012, your AWS account ID is 123456789012):

https://your\_aws\_account\_id.signin.aws.amazon.com/console/

Enter the IAM user name and password that you just created. When you're signed in, the navigation bar displays *your\_user\_name@your\_aws\_account\_id*.

If you don't want the URL for your sign-in page to contain your AWS account ID, you can create an account alias. From the IAM dashboard, click **Create Account Alias** and enter an alias, such as your company name. To sign in after you create an account alias, use the following URL:

https://your\_account\_alias.signin.aws.amazon.com/console/

To verify the sign-in link for IAM users for your account, open the IAM console and check under **AWS Account Alias** on the dashboard.

## Getting Started with Amazon Elastic File System

#### Topics

- Assumptions (p. 10)
- Related Topics (p. 11)
- Step 1: Create Your EC2 Resources and Launch Your EC2 Instance (p. 11)
- Step 2: Create Your Amazon EFS File System (p. 15)
- Step 3: Connect to Your Amazon EC2 Instance and Mount the Amazon EFS File System (p. 16)
- Step 4: Clean Up Resources and Protect Your AWS Account (p. 17)

This Getting Started exercise shows you how to quickly create an Amazon Elastic File System (Amazon EFS) file system, mount it on an Amazon Elastic Compute Cloud (Amazon EC2) instance in your VPC, and test the end-to-end setup.

There are four steps you need to perform to create and use your first Amazon EFS file system:

- Create your Amazon EC2 resources and launch your instance.
- Create your Amazon EFS file system.
- Connect to your Amazon EC2 instance and mount the Amazon EFS file system.
- Clean up your resources and protect your AWS account.

## Assumptions

For this exercise, we assume the following:

- You're already familiar with using the Amazon EC2 console to launch instances.
- Your Amazon VPC, Amazon EC2, and Amazon EFS resources are all in the same region. This guide uses the US West (Oregon) Region (us-west-2).

- You have a default VPC in the region that you're using for this Getting Started exercise. If you don't have a default VPC, or if you want to mount your file system from a new VPC with new or existing security groups, you can still use this Getting Started exercise as long as you configure Security Groups for EC2 Instances and Mount Targets (p. 69).
- You have not changed the default inbound access rule for the default security group.

You can use the root credentials of your AWS account to sign in to the console and try the Getting Started exercise. However, AWS Identity and Access Management (IAM) recommends that you do not use the root credentials of your AWS account. Instead, create an administrator user in your account and use those credentials to manage resources in your account. For more information, see Setting Up (p. 8).

## **Related Topics**

This guide also provides a walkthrough to perform a similar Getting Started exercise using AWS Command Line Interface (AWS CLI) commands to make the Amazon EFS API calls. For more information, see Walkthrough 1: Create Amazon EFS File System and Mount It on an EC2 Instance Using the AWS CLI (p. 83).

## Step 1: Create Your EC2 Resources and Launch Your EC2 Instance

Before you can launch and connect to an Amazon EC2 instance, you need to create a key pair, unless you already have one. You can create a key pair using the Amazon EC2 console and then you can launch your EC2 instance.

#### Note

Using Amazon EFS with Microsoft Windows Amazon EC2 instances is not supported.

#### To create a key pair

• Follow the steps in Setting Up with Amazon EC2 in the Amazon EC2 User Guide for Linux Instances to create a key pair. If you already have a key pair, you do not need to create a new one and you can use your existing key pair for this exercise.

#### To launch the EC2 instance

- 1. Open the Amazon EC2 console at https://console.aws.amazon.com/ec2/.
- 2. Choose Launch Instance.
- 3. In Step 1: Choose an Amazon Machine Image (AMI), find the Amazon Linux AMI at the top of the list and choose Select.

#### Note

If you choose either the Amazon Linux AMI 2016.03.0 or Amazon Linux AMI 2016.09.0 AMI when launching your Amazon EC2 instance, you won't need to install nfs-utils because it's already included in the AMI by default.

- 4. In Step 2: Choose an Instance Type, choose Next: Configure Instance Details.
- 5. In Step 3: Configure Instance Details, choose Network, and then choose the entry for your default VPC. It will look something like vpc-xxxxxx (172.31.0.0/16) (default).
  - a. Choose **Subnet**, and then choose a subnet in any Availability Zone.

#### Amazon Elastic File System User Guide Step 1: Create Your EC2 Resources and Launch Your EC2 Instance

1. Choose AMI 2. Choose Instance Type	3. Co	nfigure Instance	4. Add Storage	5. Tag Instand	ce 6. Confi	gure Security Gro	up 7. Review
Step 3: Configure Instan Configure the instance to suit your requ of the lower pricing, assign an access m	nce De irements. anagemen	etails You can launch i nt role to the inst	multiple instances tance, and more.	s from the same	e AMI, request	t Spot Instance	s to take advantage
Number of instances	1	1		]			
Purchasing option	1	Request Sp	ot Instances				
Network	(i)	vpc-	(172.31.0.0/16) (	(default)	- C	Create new \	/PC
Subnet		subnet- 4084 IP Addres	(172.31.16.0/ ses available	20)   Default in		Create new s	subnet
Auto-assign Public IP	0	Use subnet se	etting (Enable)		•		
IAM role	1	None			C C	Create new l	AM role
Shutdown behavior	1	Stop					
Enable termination protection	1	Protect again	inst accidental ter	mination			
Monitoring	1	Enable Clou Additional char	dWatch detailed	monitoring			
Tenancy	1	Shared tenand Additional char	cy (multi-tenant h rges will apply for	ardware) dedicated tena	incy.		
▼ Network interfaces ①							
Device Network Interface	Subnet	P	rimary IP	Seconda	ry IP addres	ses	
eth0 New network interface	subnet-	A	luto-assign	Add IP			
Add Device							
<ul> <li>Advanced Details</li> </ul>							
			Cancel	Previous	Review and	Launch	Next: Add Storage

- b. Choose Next: Add Storage.
- 6. Choose Next: Tag Instance.
- 7. Name your instance and choose **Next: Configure Security Group**.
- 8. In **Step 6: Configure Security Group**, review the contents of this page, ensure that **Assign a security group** is set to **Create a new security group**, and verify that the inbound rule being created has the following default values.
  - Type: SSH
  - Protocol: TCP
  - Port Range: 22
  - Source: Anywhere 0.0.0.0/0

#### Amazon Elastic File System User Guide Step 1: Create Your EC2 Resources and Launch Your EC2 Instance

1. Choose AMI	2. Choose Instance Type	3. Configure Instance	4. Add Storage	5. Tag Instance	6. Configure Security Group	7. Review
Step 6: Co A security group your instance. Fo access to the HT security groups.	onfigure Security is a set of firewall rules the or example, if you want to s TP and HTTPS ports. You	y Group at control the traffic for set up a web server and can create a new secu	your instance. On a allow Internet tra rity group or sele	this page, you car affic to reach your in ct from an existing o	n add rules to allow specific t nstance, add rules that allow one below. Learn more abo	traffic to reach v unrestricted ut Amazon EC2
	Assign a security grou	up:  Create a new	security group			
		C Select an exist	ting security grou	ıp		
	Security group nam	e: launch-wizard-4				] ,
	Descriptio	n: launch-wizard-4	created			
Туре 🕕	Prote	i)	Port Rang	je (i)	Source (i)	
SSH	TCP		22		Anywhere  0.0.0.0/	0 😣
Add Rule						
Warnir Rules w access	ng ith source of 0.0.0.0/0 allo from known IP addresses of	w all IP addresses to an	ccess your instand	ce. We recommend	setting security group rules	to allow
				Canc	el Previous Revie	w and Launch

#### Note

You can configure the EFS file system to mount on your EC2 instance automatically. For more information, see Configuring an EFS File System to Mount Automatically at EC2 Instance Launch (p. 45).

- 9. Choose Review and Launch.
- 10. Choose Launch.
- 11. Select the check box for the key pair that you created, and then choose Launch Instances.
- 12. Choose View Instances.
- 13. Choose the name of the instance you just created from the list, and then choose Actions.
  - a. From the menu that opens, choose **Networking** and then choose **Change Security Groups**.

#### Amazon Elastic File System User Guide Step 1: Create Your EC2 Resources and Launch Your EC2 Instance

🎁 AWS 🗸 Servi	ces 🗸 Edit 🗸		
EC2 Dashboard	Launch Instance Connect	Actions A	
Tags	Q. Filter by tags and attributes or sea	Connect Get Windows Password	
Limits	Name ~	Launch More Like This	Type – Availability Zone – Instance
INSTANCES	Getting Starting Example	Instance State	orunn
Instances		Image 🕨	erm 🥥 term
Spot Requests		Networking >	Change Security Groups erm
Reserved Instances		CloudWatch Monitoring >	Attach Network Interface
- IMAGES			Detach Network Interface
AMIs			Disassociate Elastic IP Address
Bundle Tasks			Change Source/Dest. Check
ELASTIC BLOCK STORE			Manage Private IP Addresses
Volumes	×		
Snapshots	Instance: (Getting St	arting Example) Public	c DNS:

- b. Select the check box next to the security group with the description **default VPC security** group.
- c. Choose Assign Security Groups.

Cha	nge Security Gro	ups		×
Insta Inter Selec	ance ID: face ID: ct Security Group(s) to a	issociate with your instance	e	
	Security Group ID	Name	Description	
		default	default VPC security group	^
				_
				-
			Cancel Assign Security	Groups

#### Note

In this step, you assign your VPC's default security group to the Amazon EC2 instance. This ensures that the instance is a member of the security group that the Amazon EFS file system mount target will authorize for connection in Step 2: Create Your Amazon EFS File System (p. 15).

By using your VPC's default security group, with its default inbound and outbound rules, you are potentially opening up this instance and this file system to potential threats from within your VPC. Make sure you follow Step 4: Clean Up Resources and Protect Your AWS Account (p. 17) at the end of this Getting Started exercise to remove resources exposed to your VPC's default security group for this example. For more information, see Security Groups for EC2 Instances and Mount Targets (p. 69).

- 14. Choose your instance from the list.
- 15. On the **Description** tab, make sure that you have two entries listed next to **security groups** one for the default VPC security group and one for the security group that you created when you launched the instance.
- 16. Make a note of the values listed next to **VPC ID** and **Public DNS**. You'll need those values later in this exercise.

## Step 2: Create Your Amazon EFS File System

In this step, you create your Amazon EFS file system.

#### To create your Amazon EFS file system

- 1. Open the Amazon EFS console at https://console.aws.amazon.com/efs/.
- 2. Choose Create File System.
- 3. Choose your default VPC from the **VPC** list. It has the same **VPC ID** that you noted at the end of Step 1: Create Your EC2 Resources and Launch Your EC2 Instance (p. 11).
- Select the check boxes for all of the Availability Zones. Make sure that they all have the default subnets, automatic IP addresses, and the default security groups chosen. These are your mount targets. For more information, see Creating Mount Targets (p. 23).

Create nie system									
Step 1: Configure file system access	Con	figure file syste	em access						
Step 2: Add tags	An An	azon EFS file system i	s accessed by EC2	instances running in	nside one of your VPCs. In	stances connect	to a file system v	a a network	
Step 3: Review and create	interfa	ice called a mount targ	et. Each mount tary	get has an IP addres	s, which we assign automa	stically or you can	specify.		
		VPC	*	0					
	Crea	te mount targets							
	Instan	ces connect to a file sy	stem via mount tar	gets you create. We	recommend creating a mo	ount target in eac	h of your VPC's A	wailability Zone	85 SO
	that E	C2 instances across yo	our VPC can acces	s the file system.					
		Availability Zone	Subnet	0	IP address	0	Security group		0
	۲		subnet-	- default *	Automatic 🥒		sg-	- default ×	
			subnet-	- default 🔹	Automatic 🥒		sg-	- default ×	
			subnet-	- default *	Automatic 🥒		50-	- default ×	
							Gancel	Next Ste	P

- 5. Choose Next Step.
- 6. Name your file system, keep **general purpose** selected as your default performance mode, and choose **Next Step**.
- 7. Choose Create File System.
- 8. Choose your file system from the list and make a note of the **File system ID** value. You'll need this value for the next step.

## Step 3: Connect to Your Amazon EC2 Instance and Mount the Amazon EFS File System

You can connect to your Amazon EC2 instance from a computer running Windows or Linux. To connect to your Amazon EC2 instance and mount the Amazon EFS file system, you need the following information:

- The **Public DNS** name of the Amazon EC2 instance. You made a note of this value at the end of Step 1: Create Your EC2 Resources and Launch Your EC2 Instance (p. 11).
- The File system ID value for the mount target for your Amazon EFS file system. You made a note of this value at the end of Step 2: Create Your Amazon EFS File System (p. 15).

#### To connect to your Amazon EC2 instance and mount the Amazon EFS file system

- 1. Connect to your Amazon EC2 instance. For more information, see Connecting to Your Linux Instance from Windows Using PuTTY or Connecting to Your Linux Instance Using SSH in the Amazon EC2 User Guide for Linux Instances.
- 2. After you've connected, install the NFS client.

If you're using an Amazon Linux AMI or RedHat Linux AMI, install the NFS client with the following command.

\$ sudo yum -y install nfs-utils

If you're using an Ubuntu AMI, install the NFS client with the following command.

\$ sudo apt-get -y install nfs-common

3. Make a directory for the mount point with the following command.

```
$ mkdir efs
```

4. Mount the Amazon EFS file system to the directory that you created. Use the following command and replace the *file-system-id* and *aws-region* placeholders with your **File System ID** value and AWS Region, respectively. Note that the command looks up your Amazon EC2 instance's Availability Zone using the instance metadata URI 169.254.169.254:

```
$ sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 $(curl
-s http://169.254.169.254/latest/meta-data/placement/availability-
zone).file-system-ID.efs.aws-region.amazonaws.com:/ efs
```

5. Change directories to the new directory that you created with the following command.

```
$ cd efs
```

6. Make a subdirectory and change the ownership of that subdirectory to your EC2 instance user. Then, navigate to that new directory with the following commands.

```
$ sudo mkdir getting-started
$ sudo chown ec2-user getting-started
$ cd getting-started
```

7. Create a text file with the following command.

```
$ touch test-file.txt
```

8. List the directory contents with the following command.

\$ ls -al

As a result, the following file is created.

```
-rw-rw-r-- 1 ec2-user ec2-user 0 Aug 15 15:32 test-file.txt
```

## Step 4: Clean Up Resources and Protect Your AWS Account

This guide includes walkthroughs that you can use to further explore Amazon EFS. Before you perform this clean-up step, you can use the resources you've created and connected to in this Getting Started exercise in those walkthroughs. For more information, see Walkthroughs (p. 83). After you have finished the walkthroughs or if you don't want to explore the walkthroughs, you should follow these steps to clean up your resources and protect your AWS account.

#### To clean up resources and protect your AWS account

- 1. Connect to your Amazon EC2 instance.
- 2. Unmount the Amazon EFS file system with the following command.

\$ sudo umount efs

- 3. Open the Amazon EFS console at https://console.aws.amazon.com/efs/.
- 4. Choose the Amazon EFS file system that you want to delete from the list of file systems.
- 5. For Actions, choose Delete file system.
- 6. In the **Permanently delete file system** dialog box, type the file system ID for the Amazon EFS file system that you want to delete, and then choose **Delete File System**.
- 7. Open the Amazon EC2 console at https://console.aws.amazon.com/ec2/.
- 8. Choose the Amazon EC2 instance that you want to terminate from the list of instances.
- 9. For Actions, choose Instance State and then choose Terminate.
- 10. In **Terminate Instances**, choose **Yes**, **Terminate** to terminate the instance that you created for this Getting Started exercise.
- 11. In the navigation pane, choose Security Groups.
- Select the name of the security group that you created for this Getting Started exercise in Step 1: Create Your EC2 Resources and Launch Your EC2 Instance (p. 11) as a part of the Amazon EC2 instance launch wizard.

#### Caution

Don't delete the default security group for your VPC.

- 13. For Actions, choose Delete Security Group.
- 14. In **Delete Security Group**, choose **Yes**, **Delete** to delete the security group you created for this Getting Started exercise.

## Creating Resources for Amazon EFS

Amazon EFS provides elastic, shared file storage that is POSIX-compliant. The file system you create supports concurrent read and write access from multiple Amazon EC2 instances and is accessible from all of the Availability Zones in the AWS Region where it is created.

You can mount an Amazon EFS file system on EC2 instances in your Amazon Virtual Private Cloud (Amazon VPC) using the Network File System version 4.1 protocol (NFSv4.1). For more information, see Amazon EFS: How It Works (p. 3).

Topics

- Creating an Amazon Elastic File System (p. 19)
- Creating Mount Targets (p. 23)
- Creating Security Groups (p. 28)

As an example, suppose you have one or more EC2 instances launched in your VPC. Now you want to create and use a file system on these instances. Following are the typical steps you need to perform to use Amazon EFS file systems in the VPC:

- Create an Amazon EFS file system When creating a file system, we recommend that you
  consider using the Name tag because the Name tag value appears in the console and makes it
  easier to identify. You can also add other optional tags to the file system.
- Create mount targets for the file system To access the file system in your VPC and mount the file system to your Amazon EC2 instance, you must create mount targets in the VPC subnets.
- **Create security groups** Both an Amazon EC2 instance and a mount target need to have associated security groups. These security groups act as a virtual firewall that controls the traffic between them. You can use the security group you associated with the mount target to control inbound traffic to your file system by adding an inbound rule to the mount target security group that allows access from a specific EC2 instance. Then, you can mount the file system only on that EC2 instance.

If you are new to Amazon EFS, we recommend that you try the following exercises that provide a firsthand, end-to-end experience of using an Amazon EFS file system:

- Getting Started (p. 10) The Getting Started exercise provides a console-based end-to-end setup in which you create a file system, mount it on an EC2 instance, and test the setup. The console takes care of many things for you and helps you set up the end-to-end experience quickly.
- Walkthrough 1: Create Amazon EFS File System and Mount It on an EC2 Instance Using the AWS CLI (p. 83) The walkthrough is similar to the Getting Started exercise, but it uses the AWS Command Line Interface (AWS CLI) to perform most of the tasks. Because the AWS CLI commands closely map to the Amazon EFS API, the walkthrough can help you familiarize yourself with the Amazon EFS API operations.

For more information about creating and accessing a file system, see the following topics.

Topics

- Creating an Amazon Elastic File System (p. 19)
- Creating Mount Targets (p. 23)
- Creating Security Groups (p. 28)

## Creating an Amazon Elastic File System

Following, you can find an explanation about how to create an Amazon EFS file system and optional tags for the file system. This section explains how to create these resources using both the console and the AWS Command Line Interface (AWS CLI).

#### Note

If you are new to Amazon EFS, we recommend you go through the Getting Started exercise, which provides console-based end-to-end instructions to create and access a file system in your VPC. For more information, see Getting Started (p. 10).

Topics

{

}

- Requirements (p. 19)
- Permissions Required (p. 20)
- Creating a File System (p. 20)

### Requirements

To create a file system, the only requirement is that you create a token to ensure idempotent operation. If you use the console, it generates the token for you. For more information, see CreateFileSystem (p. 131). After you create a file system, Amazon EFS returns the file system description as JSON. Following is an example.

```
"SizeInBytes": {
    "Value": 6144
},
"CreationToken": "console-d7f56c5f-e433-41ca-8307-9d9c0example",
"CreationTime": 1422823614.0,
"FileSystemId": "fs-c7a0456e",
"PerformanceMode" : "generalPurpose",
"NumberOfMountTargets": 0,
"LifeCycleState": "available",
"OwnerId": "231243201240"
```

If you use the console, the console displays this information in the user interface.

After creating a file system, you can create optional tags for the file system. Initially, the file system has no name. You can create a **Name** tag to assign a file system name. Amazon EFS provides the CreateTags (p. 141) operation for creating tags. Each tag is simply a key-value pair.

### Permissions Required

For all operations, such as creating a file system and creating tags, a user must have AWS Identity and Access Management permissions for the corresponding API action and resource.

You can perform any Amazon EFS operations using the root credentials of your AWS account, but using root credentials is not recommended. If you create IAM users in your account, you can grant them permissions for Amazon EFS actions with user policies. You can also use roles to grant cross-account permissions. For more information about managing permissions for the API actions, see Authentication and Access Control for Amazon EFS (p. 118).

## Creating a File System

You can create a file system using the Amazon EFS console or using the AWS Command Line Interface. You can also create file systems programmatically using AWS SDKs.

### Creating a File System Using the Amazon EFS Console

The Amazon EFS console provides an integrated experience, where you can use the user interface to specify VPC subnets to create mount targets and optional file system tags when you create a file system.

To create the file system mount targets in your VPC, you must specify VPC subnets. The console prepopulates the list of VPCs in your account that are in the selected AWS Region. First, you select your VPC, and then the console lists the Availability Zones in the VPC. For each Availability Zone, you can select a subnet from the list. After you select a subnet, you can either specify an available IP address in the subnet or let Amazon EFS choose an address.

When creating a file system, you also choose a performance mode. There are two performance modes to choose from—General Purpose and Max I/O. For the majority of use cases, we recommend that you use the general purpose performance mode for your file system. For more information about the different performance modes, see Performance Modes (p. 63).

When you choose **Create File System**, the console sends a series of API requests to create the file system, and then create tags and mount targets for the file system. The following example console shows the **MyFS** file system. It has the **Name** tag and three mount targets that are being created. The mount target lifecycle state must be **Available** before you can use it to mount the file system on an EC2 instance.

Nar	ne	File System ID	Metered Size	Number of Mount Targets	Creation D	ate
<ul> <li>Myl</li> </ul>	ExampleFileSystem	fs-8b2bce42	0.0 GiB	3	Apr 30, 201 (UTC)	5 6:59:15 PM
▼ Myl	-s	fs-f201e43b	0.0 GIB	3	May 6, 201: (UTC)	5 12:40:04 AM
	Owner ID 23124	1001-1040		Name: MyFS		
File System DNS Names EC2 mount in Mount Targe	tycle State Availa	ble			Manage Fi	le System Acces
File System DNS Names EC2 mount in Mount Targe	cycle State Availa n Access structions ts Availability Zone	ble Subnet A	Moun Idress ID	t Target Network Interface	ID Security Groups	le System Acces
File System DNS Names EC2 mount in Mount Targe	Availability Availability cone eu-west-1a	Subnet Ar subnet- 11 fbc2a39d 11	Idress ID 0.0.1.246 fsmt-c	t Target Network Interface a44a103 eni-480474	ID Security Groups e2c	Life Cycle State Creating
File System DNS Names EC2 mount in Mount Targe VPC	Availability Zone eu-west-1a eu-west-1c	Subnet IP Subnet II Rsc2a39d 11 Rsc2a39d 11 O0c36159 11	ddress Moun ID 0.0.1.246 fsml-c 0.0.19.84 fsml-c	t Target Network Interface a44a103 eni-48047r d44a104 eni-311396	ID Security Groups e2c	Life Cycle State Creating Creating

The Getting Started exercise provides console-based instructions to create and use an Amazon EFS file system. For instructions on how to create a file system using the console, see Step 1: Create Your EC2 Resources and Launch Your EC2 Instance (p. 11).

#### Creating a File System Using the AWS CLI

When using the AWS CLI, you create these resources in order. First, you create a file system. Then, you can create mount targets and optional tags for the file system using corresponding AWS CLI commands.

The following examples use the adminuser as the profile parameter value. You need to use an appropriate user profile to provide your credentials. For information about the AWS CLI, see Getting Set Up with the AWS Command Line Interface in the AWS Command Line Interface User Guide.

• To create a file system, use the Amazon EFS create-file-system CLI command (corresponding operation is CreateFileSystem (p. 131)), as shown following.

```
$ aws efs create-file-system \
--creation-token creation token \
--region aws-region \
--profile adminuser
```

For example, the following create-file-system command creates a file system in the uswest-2 region. The command specifies MyFirstFS as the creation token. For a list of AWS regions where you can create an Amazon EFS file system, see the Amazon Web Services General Reference.

```
$ aws efs create-file-system \
--creation-token MyFirstFS \
--region us-west-2 \
--profile adminuser
```

After successfully creating the file system, Amazon EFS returns the file system description as JSON, as shown in the following example.

```
{
    "SizeInBytes": {
        "Value": 6144
    },
    "CreationToken": "MyFirstFS",
    "CreationTime": 1422823614.0,
    "FileSystemId": "fs-c7a0456e",
    "PerformanceMode" : "generalPurpose",
    "NumberOfMountTargets": 0,
    "LifeCycleState": "available",
    "OwnerId": "231243201240"
}
```

Amazon EFS also provides the describe-file-systems CLI command (corresponding operation is DescribeFileSystems (p. 149)) that you can use to retrieve a list of file systems in your account, as shown following:

```
$ aws efs describe-file-systems \
--region aws-region \
--profile adminuser
```

Amazon EFS returns a list of the file systems in your AWS account created in the specified region.

• To create tags, use the Amazon EFS create-tags CLI command (the corresponding API operation is CreateTags (p. 141)). The following example command adds the Name tag to the file system.

```
aws efs create-tags \
--file-system-id File-System-ID \
--tags Key=Name,Value=SomeExampleNameValue \
--region aws-region \
--profile adminuser
```

You can retrieve a list of tags created for a file system using the describe-tags CLI command (corresponding operation is DescribeTags (p. 157)), as shown following.

```
aws efs describe-tags \
--file-system-id File-System-ID \
--region aws-region \
--profile adminuser
```

Amazon EFS returns these descriptions as JSON. The following is an example of tags returned by the DescribeTags operation. It shows a file system as having only the Name tag.

```
{
    "Tags": [
        {
            "Value": "MyFS",
            "Key": "Name"
        }
    ]
}
```

## **Creating Mount Targets**

After you create a file system, you can create mount targets and then you can mount the file system on EC2 instances in your VPC, as shown in the following illustration.



For more information about creating a file system, see Creating an Amazon Elastic File System (p. 19).

The mount target security group acts as a virtual firewall that controls the traffic. For example, it determines which Amazon EC2 instances can access the file system. This section explains the following:

- Mount target security groups and how to enable traffic.
- How to mount the file system on your Amazon EC2 instance.
- NFS-level permissions considerations.

Initially, only the root user on the Amazon EC2 instance has read-write-execute permissions on the file system. This topic discusses NFS-level permissions and provides examples that show you how to grant permissions in common scenarios. For more information, see Network File System (NFS)–Level Users, Groups, and Permissions (p. 56).

You can create mount targets for a file system using the console, using AWS Command Line Interface, or programmatically using the AWS SDKs. When using the console, you can create mount targets when you first create a file system or after the file system is created.

## Creating a Mount Target Using the Amazon EFS console

Perform the steps in the following procedure to create a mount target using the console. As you follow the console steps, you can also create one or more mount targets. You can create one mount target for each Availability Zone in your VPC.

#### To create an Amazon EFS file system (console)

- 1. Sign in to the AWS Management Console and open the Amazon EFS console at <a href="https://console.aws.amazon.com/efs/">https://console.aws.amazon.com/efs/</a>.
- 2. Choose Create File System.

	4	
	Elastic File System	
Amazon Elast	tic File System provides file storage for use instances.	with your EC2
	Create File System	
	Getting Started Guide	
		2
Create Create an EFS file system to store your files in the Amazon	Access Write files to and read files from your EFS file system via the	Nanage You can easily administer your file system, and you can view
cioud. A me system grows and sminis automatically with the files you put in, and you only pay for what you use.	versive protocol. Any number of EC2 instances can work with your file system at the same time, and your instances can be in multiple Availability Zones in a Region.	and alert on key metrics using Cloudwatch.
	Elastic File System Documentation & Support	
	Getting Started Guide   Documentation   Support   Forums	

#### Note

The console shows the preceding page only if you don't already have any Amazon EFS file systems. If you have created file systems, the console shows a list of your file systems. On the list page, choose **Create File System**.

- 3. On the **Step 1: Configure File System Access** page, select the VPC and the Availability Zone in the VPC where you want the console to create one or more mount targets for the file system that you are creating. This VPC should be the same Amazon VPC in which you created your Amazon EC2 instance in the preceding section.
  - a. Select a Amazon VPC from the VPC list.

#### Caution

If the Amazon VPC you want is not listed, verify the region in the global navigation in the Amazon EFS console.

b. In the Create Mount Targets section, select all of the Availability Zones listed.

We recommend that you create mount targets in all Availability Zones. You can then mount your file system on Amazon EC2 instances created in any of the Amazon VPC subnets.

#### Note

You can access a file system on an Amazon EC2 instance in one Availability Zone by using a mount target created in another Availability Zone, but there are costs associated with cross–Availability Zone access.

For each Availability Zone, do the following:

• Choose a **Subnet** from the list where you want to create the mount target.

You can create one mount target in each Availability Zone. If you have multiple subnets in an Availability Zone where you launched your Amazon EC2 instance, you don't have to create mount target in the same subnet, it can be any subnet in the Availability Zone.

- Leave IP Address select to Automatic. Amazon EFS will select one of the available IP addresses for the mount target.
- Specify the **Security Group** you created specifically for the mount target, or the default security group for the default VPC. Both security groups will have the necessary inbound rule that allows inbound access from the EC2 instance security group.

Click in the **Security Group** box and the console will show you the available security groups. Here you can select a specific security group and remove the **Default** security group, or leave the default in place, depending on how you configured your Amazon EC2 instance.

Create File System								
Step 1: Configure File System Access	Config	jure File Syster	n Access					
Step 2: Aug lags	An EFS file system is accessed by EC2 instances running inside one of your VPCs. Instances connect to a file system via a network interface called a mount target. Each mount target has an IP address, which we assign automatically or you can specify.							
	v	PC: vpc-460fba23	• 0					
	Create	Mount Targets						
	instances zones so	connect to a file system that EC2 instances acro	n via mount targets you creaters you creaters your VPC can access the	ate. N ne file	Ve recommend creating a r system.	nount target in e	each of your VPC's availability	
		Availability Zone	Subnet	0	IP Address	0	Security Group	3
		eu-west-1a	subnet-f8c2a39d	¥	Automatic d		× sg-2291ce47 - efs-getting-started-mt-sg	
		eu-west-1b	subnet-8c41c8fb	•	Automatic &		× sg-2291ce47 - efs-getting-started-mt-sg	
		eu-west-1c	subnet-0dc36154	¥	Automatic d <sup>e</sup>		× sg-2291ce47 - efs-getting-started-mt-sg	Ì
							Cancel Next Step	_

4. On the Step 2: Add Tags page, specify a value for the Name tag (MyExampleFileSystem).

The console prepopulates the **Name** tag because Amazon EFS uses its value as the file system display name.

Step 1: Configure File	Add Tags				
System Access Step 2: Add Tags Step 3: Review and Create	You can add tags to describe your file system. A tag consists of a case-sensitive <i>key</i> : value pair. For example, you can define a tag with <i>key</i> : value pair Corporate Department: Sales and Marketing. At a minimum, we recommend a "Name" tag.				
	Кеу	Value	Remove		
	Name	MyExampleFileSystem	Ø		
	Add New Key	Add New Value			

5. On the Step 3: Review and Create page, choose Create File System.

File System Access						
)						
rfs-getting-started-mt-sg						
efs-getting-started-mt-sg						
efs-getting-started-mt-sg						

6. The console shows the newly created file system on the **File Systems** page. Verify that all mount targets show the **Life Cycle State** as **Available**. It might take a few moments before the mount targets become available (you can expand/collapse the file system in the EFS console to force it to refresh).

You have create	d a file system. You ca	an mount your file sys	tem from a Lini	ux EC2 instance w	th an NFSv4 client install	ed. Click here for	instructions.
Create File Sys	tem Actions -						
Nam	•	File System ID	Metered	Size Numbe	er of Mount Targets	Creation Da	ite
🔹 🔹 MyE	ampleFileSystem	fs-8b2bce42	0.0 GIB	3		Apr 30, 201	5 6:59:15 PM (UTC)
Other Detail	s			Tags			Manage
				Same:	MyExampleFileSystem		
Life Cy	Owner ID cle State Availabl	e					
Life Cy File System DNS Names EC2 mount inst Mount Target	Access	e	19	Mount Target	Natwork Interface	Man	age File System Ac
Life Cy File System DNS Names EC2 mount inst Mount Target	Access Access Access Availability Zone	Subnet	IP Address	Mount Target ID	Network Interface	Man Security Groups	age File System Ac
Life Cy File System DNS Names EC2 mount inst Mount Target	Access Availability Zone eu-west-1a	Subnet subnet-18c2a39d	IP Address 10.0.1.216	Mount Target ID fsml-d017f219	Network Interface ID eni-63a8d507	Man Security Groups	age File System Ac

 Choose the DNS Names link. It shows the DNS names of the mount targets. Write down the DNS name of the mount target in the Availability Zone where you created the Amazon EC2 instance in the previous section.

In the next section, you use the DNS name to mount the file system on the Amazon EC2 instance via the mount target. The Amazon EC2 instance on which you mount the file system via the mount target can resolve the mount target DNS name to its IP address.

Now you are ready to mount the Amazon EFS file system on an Amazon EC2 instance.

### Creating a Mount Target using the AWS CLI

To create a mount target using AWS CLI, use the create-mount-target CLI command (corresponding operation is CreateMountTarget (p. 135)), as shown following.

```
$ aws efs create-mount-target \
--file-system-id file-system-id \
--subnet-id subnet-id \
--security-group ID-of-the-security-group-created-for-mount-target \
--region aws-region \
--profile adminuser
```

After successfully creating the mount target, Amazon EFS returns the mount target description as JSON as shown in the following example.

```
"MountTargetId": "fsmt-f9a14450",
"NetworkInterfaceId": "eni-385lec4e",
"FileSystemId": "fs-b6a0451f",
"LifeCycleState": "available",
"SubnetId": "subnet-b3983dc4",
"OwnerId": "23124example",
"IpAddress": "10.0.1.24"
```

{

}

You can also retrieve a list of mount targets created for a file system using the describe-mount-targets CLI command (corresponding operation is DescribeMountTargets (p. 152)), as shown following.

```
$ aws efs describe-mount-targets \
--file-system-id file-system-id \
--region aws-region \
--profile adminuser
```

For an example, see Walkthrough 1: Create Amazon EFS File System and Mount It on an EC2 Instance Using the AWS CLI (p. 83).

## **Creating Security Groups**

Both an Amazon EC2 instance and a mount target have associated security groups. These security groups act as a virtual firewall that controls the traffic between them. If you don't provide a security group when creating a mount target, Amazon EFS associates the default security group of the VPC with it.

Regardless, to enable traffic between an EC2 instance and a mount target (and thus the file system), you must configure the following rules in these security groups:

- The security groups you associate with a mount target must allow inbound access for the TCP protocol on the NFS port from all EC2 instances on which you want to mount the file system.
- Each EC2 instance that mounts the file system must have a security group that allows outbound access to the mount target on the NFS port.

For more information about security groups, see Amazon EC2 Security Groups in the Amazon EC2 User Guide for Linux Instances.

### Creating Security Groups Using the AWS Management Console

You can use the AWS Management Console to create security groups in your VPC. To connect your Amazon EFS file system to your Amazon EC2 instance, you'll need to create two security groups: one for your Amazon EC2 instance and another for your Amazon EFS mount target.

- 1. Create two security groups in your VPC. For instructions, see Creating a Security Group in the *Amazon VPC User Guide*.
- 2. In the VPC console, verify the default rules for these security groups. Both security groups should have only an outbound rule that allows traffic to leave.
- 3. You need to authorize additional access to the security groups as follows:
  - Add a rule to the EC2 security group to allow inbound access, as shown following. Optionally, you can restrict the **Source** address.

Summary	Inboun	d Rules	Outbound Rules	Tags		
Cancel Sav	/e					
Туре		Protocol		Port Range	Source	
SSH (22)		TCP (6)	Ψ.	22	0.0.0/0	0

For instructions, see Adding and Removing Rules in the Amazon VPC User Guide.

b. Add a rule to the mount target security group to allow inbound access from the EC2 security group, as shown following (where the EC2 security group is identified as the source):

Summary	Inbour	d Rules	Outbound Rules	Tags		
Cancel Save						
Туре		Protocol		Port Range	Source	1
NFS (2049)	~	TCP (6)	Ψ.	2049	sg-8b033bee	0]
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~	~~~		~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~	~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~		2

#### Note

You don't need to add an outbound rule because the default outbound rule allows all traffic to leave (otherwise, you will need to add an outbound rule to open TCP connection on the NFS port, identifying the mount target security group as the destination).

4. Verify that both security groups now authorize inbound and outbound access as described in this section.

## Creating Security Groups Using the AWS CLI

For an example that shows how to create security groups using the AWS CLI, see Step 1: Create Amazon EC2 Resources (p. 85).
## Managing Amazon EFS File Systems

File system management tasks refer to creating and deleting file systems, managing tags, and managing network accessibility of an existing file system. Managing network accessibility is about creating and managing mount targets.

You can perform these file system management tasks using the Amazon EFS console, AWS Command Line Interface (AWS CLI), or programmatically, as discussed in the following sections.

#### Topics

- Managing File System Network Accessibility (p. 30)
- Managing File System Tags (p. 37)
- Deleting an Amazon EFS File System (p. 37)

If you are new to Amazon EFS, we recommend that you try the following exercises that provide you with first-hand end-to-end experience using an Amazon EFS file system:

- Getting Started (p. 10) This exercise provides a console based end-to-end setup in which you create a file system, mount it on an EC2 instance, and test the setup. The console takes care of many things for you and thus helps you quickly set up the end-to-end experience.
- Walkthrough 1: Create Amazon EFS File System and Mount It on an EC2 Instance Using the AWS CLI (p. 83) This walkthrough is similar to the Getting Started exercise, but it uses the AWS CLI to perform most of the tasks. Because the CLI commands closely map to the Amazon EFS API, the walkthrough can help you familiarize yourself with the Amazon EFS API.

### Managing File System Network Accessibility

You mount your file system on an EC2 instance in your VPC using a mount target that you create for the file system. Managing file system network accessibility refers to managing the mount targets.



The following illustration shows how EC2 instances in a VPC access an Amazon EFS file system using a mount target.

The illustration shows three EC2 instances launched in different VPC subnets accessing an Amazon EFS file system. The illustration also shows one mount target in each of the Availability Zones (regardless of number of subnets in each Availability Zone).

You can create only one mount target per Availability Zone. If an Availability Zone has multiple subnets, as shown in one of the zones in the illustration, you create a mount target in only one of the subnets. As long as you have one mount target in an Availability Zone, the EC2 instances launched in any of its subnets can share the same mount target.

Managing mount targets refers to these activities:

• Creating and deleting mount targets in a VPC – At a minimum, you should create a mount target in each Availability Zone from which you want to access the file system.

#### Note

We recommend you create mount targets in all the Availability Zones so you can easily mount the file system on EC2 instances that you might launch in any of the Availability Zones.

If you delete a mount target, the operation forcibly breaks any mounts of the file system via the mount target being deleted, which might disrupt instances or applications using those mounts. To avoid application disruption, stop applications and unmount the file system before deleting the mount target.

You can use a file system only in one VPC at a time. That is, you can create mount targets for the file system in one VPC at a time. If you want to access the file system from another VPC, you must delete the mount targets from the current VPC and then create new mount targets in another VPC.

• Updating the mount target configuration – When you create a mount target, you associate security groups with the mount target. A security group acts as a virtual firewall that controls the traffic to and from the mount target. You can add inbound rules to controls access to the mount

target, and thus the file system. After creating a mount target, you might want to modify the security groups assigned to them.

Each mount target also has an IP address. When you create a mount target, you can choose an IP address from the subnet where you are placing the mount target. If you omit a value, Amazon EFS selects an unused IP address from that subnet.

There is no Amazon EFS operation to change the IP address after creating a mount target, so you cannot change the IP address programmatically or by using the AWS CLI. But the console enables you to change the IP address. Behind the scenes, the console deletes the mount target and creates the mount target again.

#### Caution

If you change the IP address of a mount target, you will break any existing file system mounts and you will need to remount the file system.

None of the configuration changes to file system network accessibility affect the file system itself. Your file system and data remain.

The following sections provide information about managing network accessibility of your file system.

Topics

- Creating or Deleting Mount Targets in a VPC (p. 32)
- Creating Mount Targets in Another VPC (p. 34)
- Updating the Mount Target Configuration (p. 35)

### Creating or Deleting Mount Targets in a VPC

To access an Amazon EFS file system in a VPC you need mount targets. For an Amazon EFS file system:

- You can create one mount target in each Availability Zone.
- If the VPC has multiple subnets in an Availability Zone, you can create a mount target in only one of those subnets. All EC2 instances in the Availability Zone can share the single mount target.

#### Note

We recommend that you create a mount target in each of the Availability Zones. There are cost considerations for mounting a file system on an EC2 instance in an Availability Zone through a mount target created in another Availability Zone. For more information, see Amazon EFS. In addition, by always using a mount target local to the instance's Availability Zone, you eliminate a partial failure scenario. If the mount target's zone goes down, you won't be able to access your file system through that mount target.

For more information about the operation, see CreateMountTarget (p. 135).

You can delete mount targets. Note that a mount target deletion forcibly breaks any mounts of the file system via that mount target, which might disrupt instances or applications using those mounts. For more information, see DeleteMountTarget (p. 145).

#### Using the Console

Use the following procedure to create new mount targets, delete, or update existing mount targets using the AWS Management Console.

1. In the Amazon EFS console, select the file system, choose **Actions**, and then choose **Manage File System Access**. The console displays the **Manage File System Access** page with a list of file system mount targets you have created in the selected VPC. The console shows a list of Availability Zones and mount target information, if there is a mount target in that Availability Zone.

The console shows that the file system has one mount target in the **eu-west-2c** Availability Zone, as shown following:

n EFS alled a	3 file system is acces a mount target. Each	sed by EC2 instances running mount target has an IP addres	inside one of your VPCs. Instance ss, which we assign automatically of	es connect to a file system via a network interface or you can specify.
	VPC vpc-460fba	23 🚯		
lana	ge mount target	s		
istanc o that	es connect to a file s EC2 instances acros	system via mount targets you cr ss your VPC can access the file	reate. We recommend creating a e system.	mount target in each of your VPC's availability zones
	Availability zone	Subnet	IP address	Security group
C	eu-west-1a			
c	eu-west-1b			
	eu-west-1c	subnet-00c36159	10.0.19.84	× sg-2291ce47 - efs-getting- started-mt-sg

- 2. To create new mount targets
  - a. Click on the left side in the specific Availability Zone row.
  - b. If the Availability Zone has multiple subnets, select a subnet from the Subnet list.
  - c. Amazon EFS automatically selects an available IP address, or you can provide another IP address explicitly.
  - d. Choose a Security Group from the list.

For more information about security groups, see Amazon EC2 Security Groups in the Amazon EC2 User Guide for Linux Instances.

3. To delete a mount target, choose the **X** next to the Availability Zone from which you want to remove a mount target.

#### Using the AWS CLI

To create a mount target, use the create-mount-target AWS CLI command (corresponding operation is CreateMountTarget (p. 135)), as shown following:

```
$ aws efs create-mount-target \
--file-system-id file-system-ID (for which to create the mount target) \
--subnet-id vpc-subnet-ID (in which to create mount target) \
--security-group security-group IDs (to associate with the mount target) \
--region aws-region (for example, us-west-2) \
--profile adminuser
```

Note that the AWS region (the region parameter) must be the VPC region.

You can get a list of mount targets created for a file system using the describe-mount-target AWS CLI command (corresponding operation is DescribeMountTargets (p. 152)), as shown following:

```
$ aws efs describe-mount-targets \
--file-system-id file-system-ID \
--region aws-region-where-file-system-exists \
--profile adminuser
```

Here's a sample response:

```
{
    "MountTargets": [
        {
            "MountTargetId": "fsmt-52a643fb",
            "NetworkInterfaceId": "eni-f11e8395",
            "FileSystemId": "fs-6fa144c6",
            "LifeCycleState": "available",
            "SubnetId": "subnet-15d45170",
            "OwnerId": "23124example",
            "IpAddress": "10.0.2.99"
        },
        {
            "MountTargetId": "fsmt-55a643fc",
            "NetworkInterfaceId": "eni-14a6ae4d",
            "FileSystemId": "fs-6fa144c6",
            "LifeCycleState": "available",
            "SubnetId": "subnet-0b05fc52",
            "OwnerId": "23124example",
            "IpAddress": "10.0.19.174"
        }
    ]
}
```

To delete an existing mount target, use the delete-mount-target AWS CLI command (corresponding operation is DeleteMountTarget (p. 145)), as shown following:

```
$ aws efs delete-mount-target \
--mount-target-id mount-target-ID-to-delete \
--region aws-region-where-mount-target-exists \
--profile adminuser
```

### Creating Mount Targets in Another VPC

You can use an Amazon EFS file system in one VPC at a time. That is, you create mount targets in a VPC for your file system, and use those mount targets to provide access to the file system from EC2 instances in that VPC. To access the file system from EC2 instances in another VPC, you must first delete the mount targets from the current VPC and then create new mount targets in another VPC.

#### Using the Console

1. In the Amazon EFS console, first select the file system, choose **Actions**, and then choose **Manage File System Access**.

The console displays the **Manage File System Access** page with a list of mount targets you created for the file system in a VPC. The following illustration shows a file system that has three mount targets, one in each Availability Zones.

ed a	S file system is accessed a mount target. Each mo	by EC2 instances running in unt target has an IP address	nside one of your VPCs. Instance s, which we assign automatically o	s connect to a file system via a network interface or you can specify.
	VPC vpc-460fba23	0		
ina	ege mount targets			
anc	es connect to a file syste EC2 instances across y	em via mount targets you cre our VPC can access the file	eate. We recommend creating a r system.	nount target in each of your VPC's availability zo
	Availability zone	Subnet	IP address	Security group
3	eu-west-1a	subnet-f8c2a39d	10.0.1.216	× sg-2291ce47 - efs-getting- started-mt-sg
	eu-west-1b	subnet-8c41c8fb	10.0.2.241	× sg-2291ce47 - efs-getting- started-mt-sg
3				× sg-2291ce47 - efs-getting-

2. To change the VPC, select another VPC from the **VPC** list.

The console clears all of the mount target information and lists only the Availability Zone.

- 3. Create mount targets in one or more Availability Zone as follows:
  - a. If the Availability Zone has multiple subnets, select a subnet from the **Subnet** list.
  - b. Amazon EFS automatically selects an available IP address, or you can provide another IP address explicitly.
  - c. Choose the security groups that you want to associate.

For information about security groups, see Amazon EC2 Security Groups in the Amazon EC2 User Guide for Linux Instances.

4. Choose Save.

The console first deletes the mount targets from the previous VPC and then creates new mount targets in the new VPC that you selected.

#### Using the CLI

To use a file system in another VPC, you must first delete any mount targets you previously created in a VPC and then create new mount targets in another VPC. For example AWS CLI commands, see Creating or Deleting Mount Targets in a VPC.

### Updating the Mount Target Configuration

After you create a mount target for your file system, you may want to update security groups that are in effect. You cannot change the IP address of an existing mount target. To change IP address you must delete the mount target and create a new one with the new address. Note that deleting a mount target will break any existing file system mounts.

#### Modifying the Security Group

Security groups define inbound/outbound access. When you change security groups associated with a mount target, make sure that you authorize necessary inbound/outbound access so that your EC2 instance can communicate with the file system.

For more information about security groups, see Amazon EC2 Security Groups in the Amazon EC2 User Guide for Linux Instances.

#### Using the Console

1. In the Amazon EFS console, select the file system, choose **Actions**, and then choose **Manage File System Access**.

The console displays the **Manage File System Access** page with a list of Availability Zones and mount target information, if there is a mount target in the Availability Zone.

EFS lied a	file system is access a mount target. Each n	ed by EC2 instances running nount target has an IP addres	inside one of your VPCs. Instance s, which we assign automatically o	is connect to a file system via a network interface or you can specify.
	VPC vpc-460fba23	• •		
ana	ge mount targets			
lanc that	es connect to a file sy EC2 instances across	stem via mount targets you cr your VPC can access the file	eate. We recommend creating a system.	nount target in each of your VPC's availability zone
	Availability zone	Subnet	IP address	Security group
Э	eu-west-1a	subnet-f8c2a39d	10.0.1.216	× sg-2291ce47 - efs-getting- started-mt-sg
Э	eu-west-1b	subnet-8c41c8fb	10.0.2.241	x sg-2291ce47 - efs-getting- started-mt-sg
Э	eu-west-1c	subnet-0dc36154	10.0.3.29	× sg-2291ce47 - efs-getting- started-mt-sg

 In the Security Group column, you can add or remove security groups. Choose X to remove an existing security group. Choose the Security Group box to select from other available security groups.

If you remove all security groups, Amazon EFS assigns the VPC's default security group.

#### Using the CLI

To modify security groups that are in effect for a mount target, use the modifymount-target-security-group AWS CLI command (corresponding operation is ModifyMountTargetSecurityGroups (p. 160)) to replace any existing security groups, as shown following:

```
$ aws efs modify-mount-target-security-groups \
--mount-target-id mount-target-ID-whose-configuration-to-update \
--security-groups security-group-ids-separated-by-space \
--region aws-region-where-mount-target-exists \
```

--profile adminuser

### Managing File System Tags

You can create new tags, update values of existing tags, or delete tags associated with a file system.

#### Using the Console

The console lists existing tags associated with a file system. You can add new tags, change values of existing tags, or delete existing tags.

- 1. Open the Amazon Elastic File System console at https://console.aws.amazon.com/efs/.
- 2. Choose the file system.
- 3. Choose Action and then choose Manage Tags.
- 4. On the Manage Tags page, add or delete tags. For each new tag, provide a Key and its Value.
- 5. Choose Save.

### Using the AWS CLI

You can use the create-tags CLI command to add new tags, delete-tags to delete existing tags, or use the describe-tags command to retrieve tags associated with a file system. Each CLI command corresponds to the CreateTags (p. 141), DeleteTags (p. 147), and DescribeTags (p. 157) Amazon EFS operations.

For an example walkthrough of the AWS CLI commands that you can use to add and list tags, see Step 2.1: Create Amazon EFS File System (p. 89).

The following delete-tags command removes the tag keys test1 and test2 from the tag list of the specified file system.

```
$ aws efs \
delete-tags \
--file-system-id fs-c5a1446c \
--tag-keys "test1" "test2" \
--region us-west-2 \
--profile adminuser
```

### Deleting an Amazon EFS File System

File system deletion is a destructive action that cannot be undone. You will lose the file system and any data you have in it.

#### Important

You should always unmount a file system before you delete it.

#### Using the Console

- 1. Open the Amazon Elastic File System console at https://console.aws.amazon.com/efs/.
- 2. Select the file system that you want to delete.

- 3. Choose Action and then choose Delete File System.
- 4. In **Permanently Delete File System** confirmation box, type the file system ID and then choose **Delete File System**.

The console simplifies the file deletion for you. First it deletes the associated mount targets, and then it deletes the file system.

### Using the CLI

Before you can use the AWS CLI command to delete a file system, you must delete all of the mount targets created for the file system.

For example AWS CLI commands, see Step 4: Clean up (p. 94).

### **Related Topics**

Managing Amazon EFS File Systems (p. 30)

## **Mounting File Systems**

Amazon EFS supports the Network File System version 4.1 (NFSv4.1) protocol. This section explains how to install the NFS client and how to mount your Amazon EFS file system on an Amazon EC2 instance. The section explains the mount command and the available options for specifying your file system's mount target in the mount command. The section also explains how you can use the file fstab to automatically remount your file system after any system restarts.

#### Topics

- NFS Support (p. 39)
- Installing the NFS Client (p. 40)
- Mounting on Amazon EC2 with a DNS Name (p. 41)
- Mounting with an IP Address (p. 42)
- Mounting Automatically (p. 43)
- Additional Mounting Considerations (p. 46)

### NFS Support

Amazon EFS supports the Network File System version 4.1 (NFSv4.1) and NFSv4.0 protocols when mounting your file systems on Amazon EC2 instances. While NFSv4.0 is supported, we recommend that you use NFSv4.1. Mounting your Amazon EFS file system on your Amazon EC2 instance also requires an NFS client that supports your chosen NFSv4 protocol.

In order to get the best performance out of your file system, you should use an Amazon EC2 Amazon Machine Image (AMI) that includes a Linux kernel that is version 4.0 or newer. We recommend using **Amazon Linux AMI 2016.03.0** or **Amazon Linux AMI 2016.09.0** as the AMI for the Amazon EC2 instance that you'll mount your file system to.

#### Note

Using Amazon EFS with Microsoft Windows Amazon EC2 instances is not supported.

#### **Troubleshooting AMI/Kernel Versions**

To troubleshoot issues related to certain AMI or kernel versions when using Amazon EFS from an EC2 instance, see Troubleshooting AMI and Kernel Issues (p. 81).

### Installing the NFS Client

To mount your Amazon EFS file system on your Amazon EC2 instance, first you need to install an NFS client. To connect to your EC2 instance and install an NFS client, you need the following information:

- Public DNS name of the EC2 instance and a user name to log in (ec2-user when connecting from computers running Linux or Windows).
- DNS name of your file system's mount target. You can construct this DNS name using the following generic form:

availability-zone.file-system-id.efs.aws-region.amazonaws.com

You create the mount target in the same Availability Zone where you launch your EC2 instance. For example, if you launched your EC2 instance in the us-west-2 AWS Region and in the us-west-2a Availability Zone, the mount target DNS name is:

us-west-2a.file-system-id.efs.us-west-2.amazonaws.com

The EC2 instance on which you mount the file system using the mount target can resolve the mount target DNS name to its IP address.

#### To connect your EC2 instance and install the NFS client

- 1. Connect to your EC2 instance. Note the following about connecting to the instance:
  - To connect to your instance from a computer running Mac OS or Linux, specify the .pem file to your SSH client with the -i option and the path to your private key.
  - To connect to your instance from a computer running Windows, you can use either MindTerm or PuTTY. If you plan to use PuTTY, you need to install it and use the following procedure to convert the .pem file to a .ppk file.

For more information, see the following topics in the Amazon EC2 User Guide for Linux Instances:

- Connecting to Your Linux Instance from Windows Using PuTTY
- Connecting to Your Linux Instance Using SSH

The key file cannot be publicly viewable for SSH. You can use the chmod 400 *filename*.pem command to set these permissions. For more information, see Create a Key Pair.

2. (Optional) Get updates and reboot.

```
$ sudo yum -y update
$ sudo reboot
```

After the reboot, reconnect to your EC2 instance.

3. Install the NFS client.

If you're using an Amazon Linux AMI or RedHat Linux AMI, install the NFS client with the following command.

```
$ sudo yum -y install nfs-utils
```

If you're using an Ubuntu Amazon EC2 AMI, install the NFS client with the following command.

```
$ sudo apt-get -y install nfs-common
```

If you use a custom kernel (build a custom AMI), you need to include at a minimum the NFSv4.1 client kernel module and the right NFS4 userspace mount helper.

#### Note

If you choose the **Amazon Linux AMI 2016.03.0** or **Amazon Linux AMI 2016.09.0**Amazon Linux AMI when launching your Amazon EC2 instance, you won't need to install nfs-utils because it's already included in the AMI by default.

### Mounting on Amazon EC2 with a DNS Name

There are two unique benefits to mounting your Amazon EFS file system on an Amazon EC2 instance with a DNS name. The first benefit is that the DNS name is persistent. If you delete a mount target and then create a new one in the same Availability Zone, the DNS name for that new mount target in that Availability Zone remains the same as the DNS name for the old mount target.

Second, DNS names can be generated automatically and dynamically with the script provided in this user guide and in the EFS console. Using this script to generate your DNS name prevents you from accidentally mounting your file system across the wrong Availability Zone. Since the script generates the DNS name dynamically, you can use the same script in each Availability Zone with every EFS file system in the same region.

The generic form of a mount target DNS name is:

availability-zone.file-system-id.efs.aws-region.amazonaws.com

You can get the DNS name from the console, or you can dynamically generate it. By dynamically generating your DNS name, you ensure that you are using the mount target in the same Availability Zone as your EC2 instance. In the following example, file-system-id is your EFS file system ID, and aws-region is the AWS Region in which you both launched the EC2 instance and created the file system.

```
$ sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 $(curl -s
http://169.254.169.254/latest/meta-data/placement/availability-zone).file-
system-id.efs.aws-region.amazonaws.com:/ ~/efs-mount-point
```

#### Note

The EC2 instance on which you mount the file system via the mount target must be able to resolve the DNS name to its IP address.

To be able to specify a DNS name in the mount command, the following must be true:

- The connecting EC2 instance must be inside a VPC and must be configured to use the DNS server provided by Amazon. For information about Amazon DNS server, see DHCP Options Sets in the Amazon VPC User Guide.
- The VPC of the connecting EC2 instance must have DNS host names enabled. For more information, see Viewing DNS Hostnames for Your EC2 Instance in the Amazon VPC User Guide.

If you got your DNS name from the EFS console, use a mount command without the script that dynamically generates your DNS name. After connecting to your EC2 instance and creating a

directory, efs-mount-point in this example, you can then mount the file system using the following mount command:

```
$ sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 mount-
target-DNS:/ ~/efs-mount-point
```

# Mounting on On-Premises Servers with a DNS Name

Although you can mount your file system on your on-premises server through AWS Direct Connect with a DNS name, we recommend using IP addresses for simplicity. To use DNS names, you need to integrate your DNS services in your Amazon VPC with your on-premises DNS domains. Specifically, you need to update your on-premises DNS server to forward the DNS requests for Amazon EFS mount targets to a DNS server in the Amazon VPC over the AWS Direct Connect connection. For more information, see How to Set Up DNS Resolution Between On-Premises Networks and AWS Using AWS Directory Service and Amazon Route 53, in the AWS Security Blog.

### Mounting with an IP Address

As an alternative to mounting your Amazon EFS file system with the DNS name, Amazon EC2 instances can mount a file system using a mount target's IP address. Mounting by IP address will work in environments where DNS is disabled, such as VPCs with DNS hostnames disabled, and EC2-Classic instances mounting using ClassicLink. For more information on ClassicLink, see ClassicLink in the Amazon EC2 User Guide for Linux Instances.

Mounting a file system using the mount target IP address can also be configured as a fallback option for applications configured to mount the file system using its DNS name by default. When connecting to a mount target IP address, EC2 instances should mount using the mount target IP address in the same Availability Zone as the connecting instance.

You can get the mount target IP address for your EFS file system through the console using the following procedure.

#### To obtain the mount target IP address for your EFS file system

- 1. Open the Amazon Elastic File System console at https://console.aws.amazon.com/efs/.
- 2. Choose the Name of your EFS file system from the File systems list.
- 3. In the **Mount targets** table, identify the **Availability Zone** that you want to use to mount your EFS file system to your EC2 instance.
- 4. Make a note of the IP address associated with your chosen Availability Zone.

-,								
eate file system Ad	tions -							2
Name		File system ID	N	letered size	Number of mo	unt targets	Creation date	
•								
Other details					Tags			Manage tag
					Name:			
Owner ID								
Owner ID								
Life cycle state	Available							
Life cycle state	Available							
Life cycle state	Available							
Life cycle state	Available						Manage	file system acces:
Life cycle state	Available						Manage	file system access
File system access	Available						Manage	file system acces
File system access ONS names EC2 mount instructions	Available						Manage	fie system acces
File system access ONS names EC2 mount instructions Mount targets	Available						Manage	file system acces:
File system access DNS names EC2 mount instructions Mount targets VPC	Available Availability Zone	Subnet		IP address	Mount taroet ID	Network interface ID	Manage Security groups	file system acces
File system access DNS names EC2 mount instructions Mount targets VPC	Available Availability Zone	Subnet	44.0.00	IP address	Mount target ID	Network interface ID	Manage Security groups	file system access
File system access DNS names EC2 mount instructions Mount targets VPC	Available Availability Zone	Subnet subnet-	(default)	IP address 192.0.2.9	Mount target ID	Network interface ID	Manage Security groups	file system access Life cycle state Creating
File system access DNS names EC2 mount instructions Mount targets VPC vpc- (default)	Available Availability Zone	Subnet subnet- subnet-	(defauit) (defauit)	IP address 192.0.2.0 192.0.2.1	Mount target ID	Network interface ID	Manage Security groups sg default	file system access Life cycle state Creating Available

You can specify the IP address of a mount target in the mount command, as shown following:

```
$ sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 mount-
target-IP:/ ~/efs-mount-point
```

### Mounting Automatically

You can use the file fstab to automatically mount your Amazon EFS file system whenever the Amazon EC2 instance it is mounted on reboots. There are two ways to set up automatic mounting. You can update the /etc/fstab file in your EC2 instance after you connect to the instance for the first time, or you can configure automatic mounting of your EFS file system when you create your EC2 instance.

# Updating an Existing EC2 Instance to Mount Automatically

To automatically remount your Amazon EFS file system directory when the Amazon EC2 instance reboots, you can use the file <code>fstab</code>. The file <code>fstab</code> contains information about file systems, and the command <code>mount -a</code>, which runs during instance startup, mounts the file systems listed in the <code>fstab</code> file.

#### Note

Before you can update the /etc/fstab file of your EC2 instance, make sure that you've already created your Amazon EFS file system and that you're connected to your Amazon EC2 instance. For more information, see Step 2: Create Your Amazon EFS File System (p. 15) in the Amazon EFS Getting Started exercise.

#### To update the /etc/fstab file in your EC2 instance

1. Connect to your EC2 instance, and open the /etc/fstab file in an editor.

2. Add the following line to the /etc/fstab file.

```
mount-target-DNS:/ efs-mount-point nfs4
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 0 0
```

#### Note

Do not copy the contents of your /etc/fstab file between EC2 instances in different Availability Zones. Each file system has a unique DNS name for each Availability Zone in which you have created a mount target. For more information about how to get the right DNS name for your mount point, see Mounting on Amazon EC2 with a DNS Name (p. 41).

3. Save the changes to the file.

Your EC2 instance is now configured to mount the EFS file system whenever it restarts.

#### Note

If your Amazon EC2 instance needs to start regardless of the status of your mounted Amazon EFS file system, you'll want to add the <code>nofail</code> option to your file system's entry in your <code>etc/fstab</code> file.

The line of code	you added to	the /etc/fstab	file does the	following.
------------------	--------------	----------------	---------------	------------

Field	Description
mount-target-DNS:/	The Domain Name Server (DNS) name for the file system that you want to mount. This is the same value used in mount commands to mount the subdirectory of your EFS file system.
efs-mount-point	The mount point for the EFS file system on your EC2 instance.
nfs4	The type of file system. For EFS, this type is always nfs4.
mount options	Mount options for the file system. This is a comma-separated list of the following options:
	<ul> <li>nfsvers – Identifies the version of NFS that will be used. We recommend 4.1 as the value for this option.</li> </ul>
	• rsize – Defines the size of the chunks for reading data between your client and the file system in the cloud. We recommend 1048576 as the value for this option.
	• wsize – Defines the size of the chunks for writing data between your client and the file system in the cloud. We recommend 1048576 as the value for this option.
	<ul> <li>hard – Specifies that the local applications using a file on the file system should stop and wait for the file system to come back online if Amazon EFS is temporarily unavailable.</li> </ul>
	<ul> <li>timeo – Specifies the amount of time, in tenths of a second, that the NFS client waits for a response before it retries a request to the file system in the cloud. We recommend 600 deciseconds as the value for this option.</li> </ul>
	• retrans – Specifies the number of times the NFS client should retry a request. We recommend 2 as the value for this option.
	For more information, see the section called "Additional Mounting Considerations" (p. 46).

Field	Description
0	A nonzero value indicates the file system should be backed up by ${\tt dump}.$ For EFS, this value should be 0.
0	The order in which $\tt fsck$ checks file systems at boot. For EFS file systems, this value should be 0 to indicate that $\tt fsck$ should not run at startup.

### Configuring an EFS File System to Mount Automatically at EC2 Instance Launch

You can configure an Amazon EC2 instance to mount your Amazon EFS file system automatically when it is first launched with a script that works with cloud-init. You add the script during the **Launch Instance** wizard of the EC2 management console. For an example of how to launch an EC2 instance from the console, see Getting Started (p. 10).

The script installs the NFS client and writes an entry in the /etc/fstab file that will identify the mount target DNS name as well as the subdirectory in your EC2 instance on which to mount the EFS file system. The script ensures the file gets mounted when the EC2 instance is launched and after each system reboot. Note that the mount target DNS name is dynamically constructed. For more information, see Mounting on Amazon EC2 with a DNS Name (p. 41).

For more information about the customized version of cloud-init used by Amazon Linux, see cloud-init in the Amazon EC2 User Guide for Linux Instances.

#### To configure your EC2 instance to mount an EFS file system automatically at launch

- 1. Open the Amazon EC2 console in your web browser, and begin the Launch Instance wizard.
- 2. When you reach **Step 3: Configure Instance Details**, configure your instance details, expand the **Advanced** section, and then do the following:
  - Paste the following script into **User data**. You must update the script by providing the appropriate values for *file-system-id*, *aws-region*, and *efs-mount-point*:

```
#cloud-config
package_upgrade: true
packages:
    nfs-utils
runcmd:
    mkdir -p /var/www/html/efs-mount-point/
    chown ec2-user:ec2-user /var/www/html/efs-mount-point/
    echo "$(curl -s http://169.254.169.254/latest/meta-
data/placement/availability-zone).file-system-id.efs.aws-
region.amazonaws.com:/ /var/www/html/efs-mount-point nfs4
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 0 0"
    >> /etc/fstab
    - mount -a -t nfs4
```

If you are specifying a custom path to your mount point, as in the example, you may want to use mkdir -p, because the -p option creates intermediate parent directories as needed. The - chown line of the preceding example changes the ownership of the directory at the mount point from the root user to the default Linux system user account for Amazon Linux, ec2-user. You can specify any user with this command, or leave it out of the script to keep ownership of that directory with the root user.

For more information about user data scripts, see Adding User Data in the Amazon EC2 User Guide for Linux Instances.

3. Complete the Launch Instance wizard.

#### Note

To verify that your EC2 instance is working correctly, you can integrate these steps into the Getting Started exercise. For more information, see Getting Started (p. 10).

Your EC2 instance is now configured to mount the EFS file system at launch.

### **Additional Mounting Considerations**

When mounting your Amazon EFS file system on an Amazon EC2 instance, note the following additional considerations:

• We recommend the following default Linux mount option values:

```
rsize=1048576
wsize=1048576
hard
timeo=600
retrans=2
```

- If you must change the IO size parameters (rsize and wsize), we recommend that you use the largest size possible (up to 1048576) to avoid diminished performance.
- If you must change the timeout parameter (timeo), we recommend that you use a value of at least 150, which is equivalent to 15 seconds. The timeo parameter is in deciseconds, so 15 seconds is equal to 150 deciseconds.
- We recommend that you use the hard mount option. However, if you use a soft mount, you need to set the timeo parameter to at least 150 deciseconds.
- Avoid setting any other mount options that are different from the defaults. For example, changing read or write buffer sizes, or disabling attribute caching can result in reduced performance.
- Amazon EFS ignores source ports. If you change Amazon EFS source ports, it doesn't have any effect.
- Amazon EFS does not support any of the Kerberos security variants. For example, the following will cause a mount to fail:

\$ mount -t nfs4 -o krb5p <DNS\_NAME>:/ /efs/

- When you mount an EFS file system to your Amazon EC2 instance across Availability Zones, you will incur additional network traffic costs and latency may increase. Therefore, we recommend that you keep your mount target and your EC2 instance in the same Availability Zone.
- For more mount options, and detailed explanations of the defaults, refer to the man fstab and man nfs pages.

#### **Unmounting File Systems**

Before you delete a file system, we recommend that you unmount it from every Amazon EC2 instance that it's connected to. You can unmount a file system on your Amazon EC2 instance by running the

umount command on the instance itself. You can't unmount an Amazon EFS file system through the AWS CLI, the AWS Management Console, or through any of the AWS SDKs. To unmount an Amazon EFS file system connected to an Amazon EC2 instance running Linux, use the umount command as follows:

umount ~/efs-mount-point

We recommend that you do not specify any other unount options. Avoid setting any other unount options that are different from the defaults.

You can verify that your Amazon EFS file system has been unmounted by running the df command to display the disk usage statistics for the file systems currently mounted on your Linux based Amazon EC2 instance. If the Amazon EFS file system that you want to unmounts isn't listed in the df command output, this means that the file system is unmounted.

#### Example - Identify the Mount Status of an Amazon EFS File System and Unmount It

```
$ df -T
Filesystem Type 1K-blocks Used Available Use% Mounted on
/dev/sda1 ext4 8123812 1138920 6884644 15% /
availability-zone.file-system-id.efs.aws-region.amazonaws.com :/ nfs4
9007199254740992 0 9007199254740992 0% /home/ec2-user/efs
```

```
$ umount ~/efs
```

```
$ df -T
```

Filesystem Type 1K-blocks Used Available Use% Mounted on /dev/sdal ext4 8123812 1138920 6884644 15% /

## Monitoring Amazon EFS

Monitoring is an important part of maintaining the reliability, availability, and performance of Amazon EFS and your AWS solutions. You should collect monitoring data from all of the parts of your AWS solution so that you can more easily debug a multi-point failure if one occurs. Before you start monitoring Amazon EFS, however, you should create a monitoring plan that includes answers to the following questions:

- What are your monitoring goals?
- What resources will you monitor?
- · How often will you monitor these resources?
- · What monitoring tools will you use?
- · Who will perform the monitoring tasks?
- Who should be notified when something goes wrong?

The next step is to establish a baseline for normal Amazon EFS performance in your environment, by measuring performance at various times and under different load conditions. As you monitor Amazon EFS, you should consider storing historical monitoring data. This stored data will give you a baseline to compare against with current performance data, identify normal performance patterns and performance anomalies, and devise methods to address issues.

For example, with Amazon EFS, you can monitor network throughput, I/O for read, write, and/or metadata operations, client connections, and burst credit balances for your file systems. When performance falls outside your established baseline, you might need change the size of your file system or the number of connected clients to optimize the file system for your workload.

To establish a baseline you should, at a minimum, monitor the following items:

- Your file system's network throughput.
- The number of client connections to a file system.
- The number of bytes for each file system operation, including data read, data write, and metadata operations.

### **Monitoring Tools**

AWS provides various tools that you can use to monitor Amazon EFS. You can configure some of these tools to do the monitoring for you, while some of the tools require manual intervention. We recommend that you automate monitoring tasks as much as possible.

### Automated Monitoring Tools

You can use the following automated monitoring tools to watch Amazon EFS and report when something is wrong:

- Amazon CloudWatch Alarms Watch a single metric over a time period that you specify, and perform one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The action is a notification sent to an Amazon Simple Notification Service (Amazon SNS) topic or Auto Scaling policy. CloudWatch alarms do not invoke actions simply because they are in a particular state; the state must have changed and been maintained for a specified number of periods. For more information, see Monitoring with Amazon CloudWatch (p. 49).
- Amazon CloudWatch Logs Monitor, store, and access your log files from AWS CloudTrail or other sources. For more information, see Monitoring Log Files in the Amazon CloudWatch User Guide.
- Amazon CloudWatch Events Match events and route them to one or more target functions or streams to make changes, capture state information, and take corrective action. For more information, see Using Events in the Amazon CloudWatch User Guide.
- AWS CloudTrail Log Monitoring Share log files between accounts, monitor CloudTrail log files in real time by sending them to CloudWatch Logs, write log processing applications in Java, and validate that your log files have not changed after delivery by CloudTrail. For more information, see Working with CloudTrail Log Files in the AWS CloudTrail User Guide.

### Manual Monitoring Tools

Another important part of monitoring Amazon EFS involves manually monitoring those items that the Amazon CloudWatch alarms don't cover. The Amazon EFS, CloudWatch, and other AWS console dashboards provide an at-a-glance view of the state of your AWS environment. We recommend that you also check the log files on file system.

- From the Amazon EFS console, you can find the following items for your file systems:
  - The current metered size
  - The number of mount targets
  - The life cycle state
- CloudWatch home page shows:
  - Current alarms and status
  - Graphs of alarms and resources
  - Service health status

In addition, you can use CloudWatch to do the following:

- · Create customized dashboards to monitor the services you use
- Graph metric data to troubleshoot issues and discover trends
- · Search and browse all your AWS resource metrics
- · Create and edit alarms to be notified of problems

### Monitoring with Amazon CloudWatch

You can monitor file systems using Amazon CloudWatch, which collects and processes raw data from Amazon EFS into readable, near real-time metrics. These statistics are recorded for a period of 15 months, so that you can access historical information and gain a better perspective on how

your web application or service is performing. By default, Amazon EFS metric data is automatically sent to CloudWatch at 1-minute periods. For more information about CloudWatch, see What Are Amazon CloudWatch, Amazon CloudWatch Events, and Amazon CloudWatch Logs? in the Amazon CloudWatch User Guide.

Topics

- Amazon CloudWatch Metrics for Amazon EFS (p. 50)
- Amazon EFS Dimensions (p. 52)
- How Do I Use Amazon EFS Metrics? (p. 52)
- Access CloudWatch Metrics (p. 53)
- Creating CloudWatch Alarms to Monitor Amazon EFS (p. 53)

### Amazon CloudWatch Metrics for Amazon EFS

This section includes a listing of all the available Amazon CloudWatch Metrics for Amazon EFS.

Metric	Description
BurstCreditBalance	The number of burst credits that a file system has.
	Burst credits allow a file system to burst to throughput levels above a file system's baseline level for periods of time. For more information, see Throughput scaling in Amazon EFS.
	The Minimum statistic is the smallest burst credit balance for any minute during the period. The Maximum statistic is the largest burst credit balance for any minute during the period. The Average statistic is the average burst credit balance during the period.
	Units: Bytes
	Valid statistics: Minimum, Maximum, Average
ClientConnections	The number of client connections to a file system. When using a standard client, there is one connection per mounted Amazon EC2 instance.
	Note To calculate the average ClientConnections for periods greater than one minute, divide the Sum statistic by the number of minutes in the period.
	Units: Count of client connections
	Valid statistics: Sum
DataReadIOBytes	The number of bytes for each file system read operation.
	The Sum statistic is the total number of bytes associated with read operations. The Minimum statistic is the size of the smallest read operation during the period. The Maximum statistic is the size of the largest read operation during the period. The Average statistic is the average size of read operations during the period. The SampleCount statistic provides a count of read operations.
	Units:
	• Bytes for Minimum, Maximum, Average, and Sum.
	• Count for SampleCount.

Metric	Description
	Valid statistics: Minimum, Maximum, Average, Sum, SampleCount
DataWriteIOBytes	The number of bytes for each file write operation.
	The Sum statistic is the total number of bytes associated with write operations. The Minimum statistic is the size of the smallest write operation during the period. The Maximum statistic is the size of the largest write operation during the period. The Average statistic is the average size of write operations during the period. The SampleCount statistic provides a count of write operations.
	Units:
	• Bytes are the units for the Minimum, Maximum, Average, and Sum statistics.
	• Count for SampleCount.
	Valid statistics: Minimum, Maximum, Average, Sum, SampleCount
MetadataIOBytes	The number of bytes for each metadata operation.
	The Sum statistic is the total number of bytes associated with metadata operations. The Minimum statistic is the size of the smallest metadata operation during the period. The Maximum statistic is the size of the largest metadata operation during the period. The Average statistic is the size of the size of the average metadata operation during the period. The SampleCount statistic provides a count of metadata operations.
	Units:
	<ul> <li>Bytes are the units for the Minimum, Maximum, Average, and Sum statistics.</li> <li>Count for SampleCount.</li> </ul>
	Valid statistics: Minimum, Maximum, Average, Sum, SampleCount
PercentIOLimit	Shows how close a file system is to reaching the I/O limit of the General Purpose performance mode. If this metric is at 100% more often than not, consider moving your application to a file system using the Max I/O performance mode.
	<b>Note</b> This metric is only submitted for file systems using the General Purpose performance mode.
	Units:
	• Percent

Metric	Description
PermittedThroughput	The maximum amount of throughput a file system is allowed, given the file system size and BurstCreditBalance. For more information, see Amazon EFS Performance.
	The Minimum statistic is the smallest throughput permitted for any minute during the period. The Maximum statistic is the highest throughput permitted for any minute during the period. The Average statistic is the average throughput permitted during the period.
	Units: Bytes per second
	Valid statistics: Minimum, Maximum, Average
TotalIOBytes	The number of bytes for each file system operation, including data read, data write, and metadata operations.
	The Sum statistic is the total number of bytes associated with all file system operations. The Minimum statistic is the size of the smallest operation during the period. The Maximum statistic is the size of the largest operation during the period. The Average statistic is the average size of an operation during the period. The SampleCount statistic provides a count of all operations.
	<b>Note</b> To calculate the average operations per second for a period, divide the SampleCount statistic by the number of seconds in the period. To calculate the average throughput (Bytes per second) for a period, divide the Sum statistic by the number of seconds in the period.
	Units:
	<ul> <li>Bytes for Minimum, Maximum, Average, and Sum statistics.</li> <li>Count for SampleCount.</li> </ul>
	Valid statistics: Minimum, Maximum, Average, Sum, SampleCount

### Amazon EFS Dimensions

Amazon EFS metrics use the EFS namespace and provides metrics for a single dimension, <code>FileSystemId</code>. A file system's ID can be found in the Amazon EFS management console, and it takes the form of <code>fs-xxxxxxxx</code>.

### How Do I Use Amazon EFS Metrics?

The metrics reported by Amazon EFS provide information that you can analyze in different ways. The list below shows some common uses for the metrics. These are suggestions to get you started, not a comprehensive list.

How do I?	Relevant Metrics
How can I determine my throughput?	You can monitor the daily ${\tt Sum}$ statistic of the ${\tt TotalIOBytes}$ metric to see your throughput.

How do I?	Relevant Metrics
How can I track the number of Amazon EC2 instances that are connected to a file system?	You can monitor the Sum statistic of the ClientConnections metric.
How can I see my burst credit balance?	You can see your balance by monitoring the BurstCreditBalance metric for your file system. For more information on bursting and burst credits, see Throughput Scaling in Amazon EFS (p. 64).

### Access CloudWatch Metrics

There are many ways to see the Amazon EFS metrics for CloudWatch. You can view them through the CloudWatch console, or you can access them using the CloudWatch CLI or the CloudWatch API. The following procedures show you how to access the metrics using these various tools.

#### To view metrics using the CloudWatch console

- 1. Open the CloudWatch console at https://console.aws.amazon.com/cloudwatch/.
- 2. In the navigation pane, choose **Metrics**.
- 3. Select the **EFS** namespace.
- 4. (Optional) To view a metric, type its name in the search field.
- 5. (Optional) To filter by dimension, select FileSystemId.

#### To access metrics from the AWS CLI

• Use the list-metrics command with the --namespace "AWS/EFS" namespace. For more information, see the AWS Command Line Interface Reference.

#### To access metrics from the CloudWatch API

• Call GetMetricStatistics. For more information, see Amazon CloudWatch API Reference.

#### Creating CloudWatch Alarms to Monitor Amazon EFS

You can create a CloudWatch alarm that sends an Amazon SNS message when the alarm changes state. An alarm watches a single metric over a time period you specify, and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The action is a notification sent to an Amazon SNS topic or Auto Scaling policy.

Alarms invoke actions for sustained state changes only. CloudWatch alarms will not invoke actions simply because they are in a particular state, the state must have changed and been maintained for a specified number of periods. The following procedures outlines how to create alarms for Amazon EFS.

#### Set alarms using the CloudWatch console

- 1. Sign in to the AWS Management Console and open the CloudWatch console at https:// console.aws.amazon.com/cloudwatch/.
- 2. Choose Create Alarm. This launches the Create Alarm Wizard.

- 3. Choose **EFS Metrics** and scroll through the Amazon EFS metrics to locate the metric you want to place an alarm on. To display just the Amazon EFS metrics in this dialog box, search on the file system id of your file system. Select the metric to create an alarm on and choose **Next**.
- 4. Fill in the **Name**, **Description**, **Whenever** values for the metric.
- 5. If you want CloudWatch to send you an email when the alarm state is reached, in the Whenever this alarm: field, choose State is ALARM. In the Send notification to: field, choose an existing SNS topic. If you select Create topic, you can set the name and email addresses for a new email subscription list. This list is saved and appears in the field for future alarms.

#### Note

If you use **Create topic** to create a new Amazon SNS topic, the email addresses must be verified before they receive notifications. Emails are only sent when the alarm enters an alarm state. If this alarm state change happens before the email addresses are verified, they do not receive a notification.

6. At this point, the **Alarm Preview** area gives you a chance to preview the alarm you're about to create. Choose **Create Alarm**.

#### To set an alarm using the AWS CLI

• Call put-metric-alarm. For more information, see AWS Command Line Interface Reference.

#### To set an alarm using the CloudWatch API

• Call PutMetricAlarm. For more information, see Amazon CloudWatch API Reference

## Using File Systems

After you create a file system and mount it on your EC2 instance, there are a few things you need to know in order to use it effectively:

- Users, groups, and related NFS-Level permissions management When you first create the file system, there is only one root directory at /. By default, only the root user (UID 0) has read-write-execute permissions. In order for other users to modify the file system, the root user must explicitly grant them access. For more information, see Network File System (NFS)–Level Users, Groups, and Permissions (p. 56).
- **Metering** Amazon EFS reports file system sizes and sizes of objects within a file system. You can view space usage in the console. For more information about how Amazon EFS reports the file system sizes, see Metering How Amazon EFS Reports File System and Object Sizes (p. 58).
- Unsupported NFSv4 features Amazon EFS supports NFSv4, however some of the NFSv4 features are not supported. For more information, see Unsupported NFSv4 Features (p. 60).

### Data Consistency in Amazon EFS

Amazon EFS provides the open-after-close consistency semantics that applications expect from NFS.

In Amazon EFS, write operations will be durably stored across Availability Zones when:

- An application performs a synchronous write operation (for example, using the open Linux command with the O\_DIRECT flag, or the fsync Linux command).
- · An application closes a file.

Amazon EFS provides stronger consistency guarantees than open-after-close semantics depending on the access pattern. Applications that perform synchronous data access and perform non-appending writes will have read-after-write consistency for data access.

### **Related Topics**

Amazon EFS: How It Works (p. 3)

Getting Started (p. 10)

Walkthroughs (p. 83)

# Network File System (NFS)–Level Users, Groups, and Permissions

#### Topics

- Example Amazon EFS File System Use Cases and Permissions (p. 56)
- User and group ID permissions on files and directories within a file system (p. 57)
- No Root Squashing (p. 58)
- Permissions Caching (p. 58)
- Changing File System Object Ownership (p. 58)

After creating a file system, by default, only the root user (UID 0) has read-write-execute permissions. In order for other users to modify the file system, the root user must explicitly grant them access.

Amazon EFS file system objects have a Unix-style mode associated with them. This value defines the permissions for performing actions on that object, and users familiar with Unix-style systems can easily understand how Amazon EFS behaves with respect to these permissions.

Additionally, on Unix-style systems, users and groups are mapped to numeric identifiers, which Amazon EFS uses to represent file ownership. File system objects (that is, files, directories, etc.) on Amazon EFS are owned by a single owner and a single group. Amazon EFS uses these numeric IDs to check permissions when a user attempts to access a file system object.

This section provides examples of permissions and discusses Amazon EFS–specific NFS permissions considerations.

# Example Amazon EFS File System Use Cases and Permissions

After you create an Amazon EFS file system and mount targets for the file system in your VPC, you can mount the remote file system locally on your Amazon EC2 instance. The mount command can mount any directory in the file system. However, when you first create the file system, there is only one root directory at /.

The following mount command mounts the root directory of an Amazon EFS file system, identified by the mount target IP address, on the /efs-mount-point local directory.

```
sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 mount-
target-ip:/ /efs-mount-point
```

Note that the root user and root group own the mounted directory.

\$ ls -al									
total 8									
drwxr-xr-x	2	root	root	4096	Jan	30	20:55		
drwx	4	ec2-user	ec2-user	4096	Feb	3	23:49	• •	

The initial permissions mode allows:

• read-write-execute permissions to the owner root

- read-execute permissions to the group root
- read-execute permissions to others

Note that only the root user can modify this directory. The root user can also grant other users permissions to write to this directory. For example:

- Create writable per-user subdirectories. For step-by-step instructions, see Walkthrough 3: Create Writable Per-User Subdirectories and Configure Automatic Remounting on Reboot (p. 101).
- Allow users to write to the Amazon EFS file system root. A user with root privileges can grant other users access to the file system.
  - To change the Amazon EFS file system ownership to a non-root user and group, use the following:

\$ sudo chown user:group /EFSroot

• To change permissions of the file system to something more permissive, use the following:

\$ sudo chmod 777 / EFSroot

This command grants read-write-execute privileges to all users on all EC2 instances that have the file system mounted.

# User and group ID permissions on files and directories within a file system

Files and directories in an Amazon EFS file system support standard Unix-style read/write/execute permissions based on the user ID and group ID asserted by the mounting NFSv4.1 client. When a user attempts to access files and directories, Amazon EFS checks their user ID and group IDs to verify the user has permission to access the objects. Amazon EFS also uses these IDs as the owner and group owner for new files and directories the user creates. Amazon EFS does not examine user or group names—it only uses the numeric identifiers.

#### Note

When you create a user on an EC2 instance, you can assign any numeric UID and GID to the user. The numeric user IDs are set in the /etc/passwd file on Linux systems. The numeric group IDs are in the /etc/group file. These files define the mappings between names and IDs. Outside of the EC2 instance, Amazon EFS does not perform any authentication of these IDs, including the root ID of 0.

If a user accesses an Amazon EFS file system from two different EC2 instances, depending on whether the UID for the user is the same or different on those instances, you see different behavior as follows:

- If the user IDs are the same on both EC2 instances, Amazon EFS considers them to be the same user, regardless of the EC2 instance they use. The user experience when accessing the file system is the same from both EC2 instances.
- If the user IDs are not the same on both EC2 instances, Amazon EFS considers them to be different users, and the user experience will not be the same when accessing the Amazon EFS file system from the two different EC2 instances.
- If two different users on different EC2 instances share an ID, Amazon EFS considers them the same user.

You might consider managing user ID mappings across EC2 instances consistently. Users can check their numeric ID using the *id* command, as shown following:

\$ id

uid=502(joe) gid=502(joe) groups=502(joe)

#### Turn Off the ID Mapper

The NFS utilities in the operating system include a daemon called an ID Mapper that manages mapping between user names and IDs. In Amazon Linux, the daemon is called rpc.idmapd and on Ubuntu is called idmapd. It translates user and group IDs into names, and vice versa. However, Amazon EFS deals only with numeric IDs. We recommend you turn this process off on your EC2 instances (on Amazon Linux the mapper is usually disabled, in which case don't enable the ID mapper), as shown following:

\$ service rpcidmapd status
\$ sudo service rpcidmapd stop

### No Root Squashing

When root squashing is enabled, the root user is converted to a user with limited permissions on the NFS server.

Amazon EFS behaves like a Linux NFS server with  $no\_root\_squash$ . If a user or group ID is 0, Amazon EFS treats that user as the root user, and bypasses permissions checks (allowing access and modification to all file system objects).

#### **Permissions Caching**

Amazon EFS caches file permissions for a small time period. As a result, there may be a brief window where a user who had access to a file system object but the access was revoked recently can still access that object.

### Changing File System Object Ownership

Amazon EFS enforces the POSIX chown\_restricted attribute. This means only the root user can change the owner of a file system object. While the root or the owner user can change the owner group of a file system object, unless the user is root, the group can only be changed to one that the owner user is a member of.

### Metering – How Amazon EFS Reports File System and Object Sizes

This section explains how Amazon EFS reports file system sizes and sizes of objects within a file system.

### Metering Amazon EFS File System Objects

Customer-visible objects in an Amazon EFS system can be regular files, directories, symbolic links, and special files (FIFOs and sockets). Each of these objects is metered for 2 KiB (kibibytes) of metadata (for its inode) and one or more increments of 4 KiB of data. The following list explains the metered data size for different types of file system objects.

• **Regular files** – The metered data size of a regular file is the logical size of the file rounded to the next 4 KiB increment, except that it may be less for sparse files.

A sparse file is a file to which data is not written to all positions of the file before its logical size is reached. For a sparse file, if the actual storage used is less than the logical size rounded to the next 4 KiB increment, Amazon EFS reports actual storage used as the metered data size.

- **Directories** The metered data size of a directory is the actual storage used for the directory entries and the data structure that holds them, rounded to the next 4 KiB increment (it does not include the actual storage used by the file data).
- Symbolic links and special files The metered data size for these objects is always 4 KiB.

When Amazon EFS reports the space used for an object, through the NFSv4.1  $space\_used$  attribute, it includes the object's current metered data size, but not its metadata size. There are two utilities available for measuring the disk usage of a file, the du and stat utilities. Here's an example of how to use the du utility, on an empty file, with the -k option to return the output in kilobytes:

```
$ du -k file
4 file
```

Here's an example of how to use the stat utility on an empty file to return the file's disk usage:

```
$ /usr/bin/stat --format="%b*%B" file | bc
4096
```

To measure the size of a directory, use the stat utility, find the Blocks value, and then multiply that value by the block size. Here's an example of how to use the stat utility on an empty directory:

```
$ /usr/bin/stat --format="%b*%B" . | bc
4096
```

#### Metering an Amazon EFS File System

The metered size of an entire Amazon EFS file system is the sum of the sizes (including metadata) of all of its current objects. The size of each object is calculated from a representative sampling that represents the size of the object during the metered hour, for example the hour from 8:00 am to 9:00 am.

For example, an empty file contributes 6 KiB (2 KiB metadata + 4 KiB data) to the metered size of its file system. Upon creation, a file system has a single empty root directory and therefore has a metered size of 6 KiB.

The metered sizes of a particular file system define the usage for which the owner account is billed for that file system for that hour.

#### Note

The computed metered size does not represent a consistent snapshot of the file system at any particular time during that hour. Rather, it represents the sizes of the objects that existed in the file system at varying times within each hour or possibly the hour before it, which are summed to determine the file system's metered size for the hour. The metered size of a file systems is thus eventually consistent with the metered sizes of the objects stored when there are no writes to the file system.

This metered size for an Amazon EFS file system can be seen in the following ways:

• DescribeFileSystems API – Used in SDKs, HTTP, and the AWS CLI.

- File Systems table For each file system listed in the AWS Management Console.
- DF command In Linux, the df command can be run at the terminal prompt of an EC2 instance.

#### Note

The metered size is also used to determine your I/O throughput baseline and burst rates. For more information, see Throughput Scaling in Amazon EFS (p. 64).

### **Unsupported NFSv4 Features**

While Amazon Elastic File System does not support NFSv2, or NFSv3, Amazon EFS supports both NFSv4.1 and NFSv4.0, except for the following features:

- pNFS
- Client delegation or callbacks of any type
  - Operation OPEN always returns OPEN\_DELEGATE\_NONE as the delegation type.
  - The operation OPEN returns NFSERR\_NOTSUPP for the CLAIM\_DELEGATE\_CUR and CLAIM\_DELEGATE\_PREV claim types.
- Mandatory locking

All locks in Amazon EFS are advisory, which means that READ and WRITE operations do not check for conflicting locks before the operation is executed.

· Deny share

NFS supports the concept of a share deny, primarily used by Windows clients for users to deny others access to a particular file that has been opened. Amazon EFS does not support this, and returns the NFS error NFS4ERR\_NOTSUPP for any OPEN commands specifying a share deny value other than OPEN4\_SHARE\_DENY\_NONE. Linux NFS clients do not use anything other than OPEN4\_SHARE\_DENY\_NONE.

- Access control lists (ACL)
- Amazon EFS does not update the time\_access attribute on file reads. Amazon EFS updates time\_access in the following events:
  - When a file is created (an inode is created).
  - When NFS client makes an explicit setattr call.
  - On a write to the inode caused by, for example, file size changes or file metadata changes.
  - Any inode attribute is updated.
- Namespaces
- Persistent reply cache
- · Kerberos based security
- NFSv4.1 data retention
- Lock upgrades and downgrades: Amazon EFS returns NFS4ERR\_LOCK\_NOTSUPP if the client attempts to upgrade or downgrade an existing lock.

#### Note

Because lock upgrades and downgrades are not supported, use cases that require this feature, such as those using SQLite or IPython, are also not supported in Amazon EFS.

- Both setuid and setgid are not supported in Amazon EFS.
- Unsupported file types when using the CREATE operation: Block devices (NF4BLK), character devices (NF4CHR), attribute directory (NF4ATTRDIR), and named attribute (NF4NAMEDATTR).
- Unsupported attributes: FATTR4\_ARCHIVE, FATTR4\_FILES\_AVAIL, FATTR4\_FILES\_FREE, FATTR4\_FILES\_TOTAL, FATTR4\_FS\_LOCATIONS, FATTR4\_MIMETYPE,

FATTR4\_QUOTA\_AVAIL\_HARD, FATTR4\_QUOTA\_AVAIL\_SOFT, FATTR4\_QUOTA\_USED, FATTR4\_TIME\_BACKUP, and FATTR4\_ACL.

An attempt to set these attributes will result in an NFS4ERR\_ATTRNOTSUPP error that is sent back to the client.

## Amazon EFS Performance

This topic provides an overview of Amazon EFS performance, discusses the two performance modes (General Purpose and Max I/O) available in Amazon EFS, reviews the Amazon EFS bursting model, and outlines some useful performance tips.

### Performance Overview

Amazon EFS file systems are distributed across an unconstrained number of storage servers, enabling file systems to grow elastically to petabyte scale and allowing massively parallel access from Amazon EC2 instances to your data. Amazon EFS's distributed design avoids the bottlenecks and constraints inherent to traditional file servers.

This distributed data storage design means that multithreaded applications and applications that concurrently access data from multiple Amazon EC2 instances can drive substantial levels of aggregate throughput and IOPS. Big data and analytics workloads, media processing workflows, content management, and web serving are examples of these applications.

In addition, Amazon EFS data is distributed across multiple Availability Zones (AZs), providing a high level of durability and availability. The following tables compare high-level performance and storage characteristics for Amazon's file and block cloud storage services.

	Amazon EFS	Amazon EBS Provisioned IOPS		
Per-operation latency	Low, consistent latency.	Lowest, consistent latency.		
Throughput scale	Multiple GBs per second.	About 1 GB per second.		

#### Performance Comparison – Amazon EFS and Amazon EBS

#### Storage Characteristics Comparison – Amazon EFS and Amazon EBS

	Amazon EFS	Amazon EBS Provisioned IOPS
Availability and durability	Data is stored redundantly across multiple AZs.	Data is stored redundantly in a single AZ.
Access	Up to thousands of Amazon EC2 instances, from multiple AZs, can connect concurrently to a file system.	A single Amazon EC2 instance in a single AZ can connect to a file system.

	Amazon EFS	Amazon EBS Provisioned IOPS		
Use cases	Big data and analytics, media processing workflows, content management, web serving, and home directories.	Boot volumes, transactional and NoSQL databases, data warehousing, and ETL.		

The distributed nature of Amazon EFS enables high levels of availability, durability, and scalability. This distributed architecture results in a small latency overhead for each file operation. Due to this peroperation latency, overall throughput generally increases as the average I/O size increases, because the overhead is amortized over a larger amount of data. Amazon EFS supports highly parallelized workloads (for example, using concurrent operations from multiple threads and multiple Amazon EC2 instances), which enables high levels of aggregate throughput and operations per second.

### Amazon EFS Use Cases

Amazon EFS is designed to meet the performance needs of the following use cases.

### Big Data and Analytics

Amazon EFS provides the scale and performance required for big data applications that require high throughput to compute nodes coupled with read-after-write consistency and low-latency file operations.

### Media Processing Workflows

Media workflows like video editing, studio production, broadcast processing, sound design, and rendering often depend on shared storage to manipulate large files. A strong data consistency model with high throughput and shared file access can cut the time it takes to perform these jobs and consolidate multiple local file repositories into a single location for all users.

### Content Management and Web Serving

Amazon EFS provides a durable, high throughput file system for content management systems that store and serve information for a range of applications like websites, online publications, and archives.

### Home Directories

Amazon EFS can provide storage for organizations that have many users that need to access and share common data sets. An administrator can use Amazon EFS to create a file system accessible to people across an organization and establish permissions for users and groups at the file or directory level.

### **Performance Modes**

To support a wide variety of cloud storage workloads, Amazon EFS offers two performance modes. You select a file system's performance mode when you create it.

The two performance modes have no additional costs, so your Amazon EFS file system is billed and metered the same, regardless of your performance mode. For information about file system limits, see Limits for Amazon EFS File Systems (p. 73).

#### Note

An Amazon EFS file system's performance mode can't be changed after the file system has been created.

### General Purpose Performance Mode

We recommend the General Purpose performance mode for the majority of your Amazon EFS file systems. General Purpose is ideal for latency-sensitive use cases, like web serving environments, content management systems, home directories, and general file serving. If you don't choose a performance mode when you create your file system, Amazon EFS selects the General Purpose mode for you by default.

### Max I/O Performance Mode

File systems in the Max I/O mode can scale to higher levels of aggregate throughput and operations per second with a tradeoff of slightly higher latencies for file operations. Highly parallelized applications and workloads, such as big data analysis, media processing, and genomics analysis, can benefit from this mode.

### Using the Right Performance Mode

Our recommendation for determining which performance mode to use is as follows:

- 1. Create a new file system (p. 15) using the default General Purpose performance mode.
- 2. Run your application (or a use case similar to your application) for a period of time to test its performance.
- 3. Monitor the PercentIOLimit (p. 50) Amazon CloudWatch metric for Amazon EFS during the performance test. For more information about accessing this and other metrics, see Amazon CloudWatch Metrics (p. 48).

If the PercentIOLimit percentage returned was at or near 100% for a significant amount of time during the test, your application should use the Max I/O performance mode. Otherwise, it should use the default General Purpose mode.

### Throughput Scaling in Amazon EFS

Throughput on Amazon EFS scales as a file system grows. Because file-based workloads are typically spiky—driving high levels of throughput for short periods of time, and low levels of throughput the rest of the time—Amazon EFS is designed to burst to high throughput levels for periods of time.

All file systems, regardless of size, can burst to 100 MiB/s of throughput, and those over 1 TiB large can burst to 100 MiB/s per TiB of data stored in the file system. For example, a 10 TiB file system can burst to 1,000 MiB/s of throughput (10 TiB  $\times$  100 MiB/s/TiB). The portion of time a file system can burst is determined by its size, and the bursting model is designed so that typical file system workloads will be able to burst virtually any time they need to.

Amazon EFS uses a credit system to determine when file systems can burst. Each file system earns credits over time at a baseline rate that is determined by the size of the file system, and uses credits whenever it reads or writes data. The baseline rate is 50 MiB/s per TiB of storage (equivalently, 50 KiB/ s per GiB of storage).

Accumulated burst credits give the file system permission to drive throughput above its baseline rate. A file system can drive throughput continuously at its baseline rate, and whenever it's inactive or driving throughput below its baseline rate, the file system accumulates burst credits.

For example, a 100 GiB file system can burst (at 100 MiB/s) for 5% of the time if it's inactive for the remaining 95%. Over a 24-hour period, the file system earns 432,000 MiBs worth of credit, which can be used to burst at 100 MiB/s for 72 minutes.

File systems larger than 1TiB can always burst for up to 50% of the time if they are inactive for the remaining 50%.

The following table provides examples of bursting behavior.

File System Size	Aggregate Read/Write Throughput
A 100 GiB file system can	<ul><li>Burst to 100 MiB/s for up to 72 minutes each day, or</li><li>Drive up to 5 MiB/s continuously</li></ul>
A 1 TiB file system can	<ul><li>Burst to 100 MiB/s for 12 hours each day, or</li><li>Drive 50 MiB/s continuously</li></ul>
A 10 TiB file system can	<ul><li>Burst to 1 GiB/s for 12 hours each day, or</li><li>Drive 500 MiB/s continuously</li></ul>
Generally, a larger file system can	<ul> <li>Burst to 100MiB/s per TiB of storage for 12 hours each day, or</li> <li>Drive 50 MiB/s per TiB of storage continuously</li> </ul>

#### Note

- The minimum file system size used when calculating the baseline rate is 1 GiB, so all file systems have a baseline rate of at least 50 KiB/s.
- The file system size used when determining the baseline rate and burst rate is the same as the metered size available through the DescribeFileSystems operation.
- File systems can earn credits up to a maximum credit balance of 2.1 TiB for file systems smaller than 1 TiB, or 2.1 TiB per TiB stored for file systems larger than 1 TiB. This implies that file systems can accumulate enough credits to burst for up to 12 hours continuously.
- Newly created file systems begin with an initial credit balance of 2.1 TiB, which enables them to add data at the 100 MiB/s burst rate until they are large enough to run at 100 MiB/s continuously (that is, 2 TiB).

The following table provides more detailed examples of bursting behavior for file systems of different sizes.

File System Size (GiB)	Baseline Aggregate Throughput (MiB/s)	Burst Aggregate Throughput (MiB/s)	Maximum Burst Duration (Min/ Day)	% of Time File System Can Burst (Per Day)
10	0.5	100	7.2	0.5%
256	12.5	100	180	12.5%
512	25.0	100	360	25.0%
1024	50.0	100	720	50.0%
1536	75.0	150	720	50.0%
2048	100.0	200	720	50.0%
File System Size (GiB)	Baseline Aggregate Throughput (MiB/s)	Burst Aggregate Throughput (MiB/s)	Maximum Burst Duration (Min/ Day)	% of Time File System Can Burst (Per Day)
---------------------------	------------------------------------------------	------------------------------------------	-----------------------------------------	-------------------------------------------------
3072	150.0	300	720	50.0%
4096	200.0	400	720	50.0%

### Note

As previously mentioned, new file systems have an initial burst credit balance of 2.1 TB. With this starting balance, you can burst at 100 MB/s for 6.12 hours (which is calculated by 2.1  $\times$  1024  $\times$  (1024/100/3600) to get 6.116 hours, rounded up to 6.12) without spending any credits that you're earning from your storage.

## Managing Burst Credits

When a file system has a positive burst credit balance, it can burst. You can see the burst credit balance for a file system by viewing the BurstCreditBalance Amazon CloudWatch metric for Amazon EFS. For more information about accessing this and other metrics, see Monitoring Amazon EFS (p. 48).

The bursting capability (both in terms of length of time and burst rate) of a file system is directly related to its size. Larger file systems can burst at larger rates for longer periods of time. Therefore, if your application needs to burst more (that is, if you find that your file system is running out of burst credits), you should increase the size of your file system.

### Note

There's no provisioning with Amazon EFS, so to make your file system larger you need to add more data to it.

Use your historical throughput patterns to calculate the file system size you need to sustain your desired level of activity. The following steps outline how to do this:

- 1. Identify your throughput needs by looking at your historical usage. From the Amazon CloudWatch console, check the sum statistic of the TotalIOBytes metric with daily aggregation, for the past 14 days. Identify the day with the largest value for TotalIOBytes.
- 2. Divide this number by 24 hours, 60 minutes, 60 seconds, and 1024 bytes to get the average KiB/ second your application required for that day.
- Calculate the file system size (in GB) required to sustain this average throughput by dividing the average throughput number (in KB/s) by the baseline throughput number (50 KB/s/GiB) that EFS provides.

# **On-Premises Performance Considerations**

The throughput bursting model for Amazon EFS file systems remains the same whether accessed from your on-premises servers or your Amazon EC2 instances. However, when accessing Amazon EFS file data from your on-premises servers, the maximum throughput is also constrained by the bandwidth of the AWS Direct Connect connection.

Because of the propagation delay tied to data traveling over long distances, the network latency of an AWS Direct Connect connection between your on-premises datacenter and your Amazon VPC can be tens of milliseconds. If your file operations are serialized, the latency of the AWS Direct Connect connection directly impacts your read and write throughput. In essence, the volume of data you can

read or write during a period of time is bounded by the amount of time it takes for each read and write operation to complete. To maximize your throughput, parallelize your file operations so that multiple reads and writes are processed by Amazon EFS concurrently. Standard tools like GNU parallel enable you to parallelize the copying of file data.

## Architecting for High Availability

To ensure continuous availability between your on-premises data center and your Amazon VPC, we recommend configuring two AWS Direct Connect connections. For more information, see Step 4: Configure Redundant Connections with AWS Direct Connect in the AWS Direct Connect User Guide.

To ensure continuous availability between your application and Amazon EFS, we recommend that your application be designed to recover from potential connection interruptions. In general, there are two scenarios for on-premises applications connected to an Amazon EFS file system; highly available and not highly available.

If your application is Highly Available (HA) and uses multiple on-premises servers in its HA cluster, ensure that each on-premises server in the HA cluster connects to a mount target in a different Availability Zone (AZ) in your Amazon VPC. If your on-premises server can't access the mount target because the AZ in which the mount target exists becomes unavailable, your application should failover to a server with an available mount target.

If your application is not highly available, and your on-premises server can't access the mount target because the AZ in which the mount target exists becomes unavailable, your application should implement restart logic and connect to a mount target in a different AZ.

# Amazon EFS Performance Tips

When using Amazon EFS, keep the following performance tips in mind:

- Average I/O Size Amazon EFS's distributed nature enables high levels of availability, durability, and scalability. This distributed architecture results in a small latency overhead for each file operation. Due to this per-operation latency, overall throughput generally increases as the average I/ O size increases, because the overhead is amortized over a larger amount of data.
- **Simultaneous Connections** Amazon EFS file systems can be mounted on up to thousands of Amazon EC2 instances concurrently. If your application is parallelizable across more instances, you can drive higher throughput levels on your file system in aggregate across instances.
- **Request Model** By enabling asynchronous writes to your file system, pending write operations are buffered on the Amazon EC2 instance before they are written to Amazon EFS asynchronously. Asynchronous writes typically have lower latencies. When performing asynchronous writes, the kernel uses additional memory for caching. A file system that has enabled synchronous writes, or one that opens files using an option that bypasses the cache (for example, O\_DIRECT), will issue synchronous requests to Amazon EFS and every operation will go through a round trip between the client and Amazon EFS.

### Note

Your chosen request model will have tradeoffs in consistency (if you're using multiple Amazon EC2 instances) and speed.

• NFS Client Mount Settings – Verify that you're using the recommended mount options as outlined in Mounting File Systems (p. 39) and in Additional Mounting Considerations (p. 46). Amazon EFS supports the Network File System version 4.1 (NFSv4.1) and NFSv4.0 protocols when mounting your file systems on Amazon EC2 instances. NFSv4.1 provides better performance.

### Note

You might want to increase the size of the read and write buffers for your NFS client to 1 MB when you mount your file system.

• Amazon EC2 Instances – Applications that perform a large number of read and write operations likely need more memory or computing capacity than applications that don't. When launching your Amazon EC2 instances, choose instance types that have the amount of these resources that your application needs. Note that the performance characteristics of Amazon EFS file systems are not dependent on the use of EBS-optimized instances.

For information about the Amazon EFS limits for total file system throughput, per-instance throughput, and operations per second in General Purpose performance mode, see Amazon EFS Limits (p. 72).

# **Related Topics**

- Metering How Amazon EFS Reports File System and Object Sizes (p. 58)
- Troubleshooting Amazon EFS (p. 75)

# Security

This section describes security considerations for working with Amazon EFS. There are three levels of access control to consider for Amazon EFS file systems, with different mechanisms used for each.

### Topics

- AWS Identity and Access Management (IAM) permissions for API calls (p. 69)
- Security Groups for EC2 Instances and Mount Targets (p. 69)
- Read, Write, and Execute Permissions for Files and Directories (p. 71)

# AWS Identity and Access Management (IAM) permissions for API calls

Example: IAM user Alice has permissions to retrieve descriptions of all file systems in her parent AWS account, but she is only allowed to manage the security groups for one of them, file system ID fs-12345678.

File systems are created, managed, and deleted with calls to the Amazon EFS API. If the caller is using credentials for an IAM user or assumed role, each API call requires that the caller have permissions for the action being called in its IAM policy. Some API actions support policy permissions specific to the file system that is the object of the call (that is, resource-level permissions). API calls made with an account's root credentials have permissions for all API actions on file systems owned by the account. For more information about IAM permissions with the Amazon EFS API, see Authentication and Access Control for Amazon EFS (p. 118).

# Security Groups for EC2 Instances and Mount Targets

When using Amazon EFS, you associate one Amazon EC2 security group with an EC2 instance and another with the EFS mount target associated with the file system. Security groups act as a firewall, and the rules you add define the traffic flow. In the Getting Started exercise, you created one security group when you launched the EC2 instance, and you associated another with the EFS mount target (that is, the default security group for your default VPC). That is fine for the Getting Started exercise, but you should set up security groups with minimal permissions for use with EFS.

You can authorize inbound/outbound access to your EFS file system by adding rules that allow your EC2 instance to connect to your Amazon EFS file system via the mount target using the NFS port. Follow these steps to create and update your security groups.

### To create your security groups

1. Create two security groups in your VPC.

For instructions, see Creating a Security Group in the Amazon VPC User Guide, follow the steps in the *To create a security group* procedure.

2. In the VPC console, verify the default rules for these security groups. Both security groups should have only an outbound rule that allows traffic to leave.

### To update the necessary access for your security groups

For the EC2 security group, add a rule to allow inbound access using SSH from any host. If you
want, you can optionally restrict the **Source** address. You don't need to add any outbound rule
because the default outbound rule allows all traffic to leave (otherwise, you'd need to add an
outbound rule to open the TCP connection on the NFS port, identifying the mount target security
group as the destination).

Summary	Inbou	nd Rules	Outbound Rules	Tags		
Cancel Sa	ive					
Гуре		Protocol		Port Range	Source	
SSH (22)		TCP (6)	Ŧ	22	0.0.0/0	0

For instructions, see Adding and Removing Rules in the Amazon VPC User Guide.

2. For the mount target security group, add a rule to allow inbound access from the EC2 security group as shown (the EC2 security group is identified as the source).

Summary	Inbound Rules	Outbound Rules	Tags		
Cancel Save	2				
Туре	Protocol		Port Range	Source	100
NFS (2049)	TCP (6)	V	2049	sg-8b033bee	0]
~	~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~	~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~	~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~		در

3. Verify that both security groups now authorize inbound and outbound access as described in this procedure.

For more information about security groups, see Security Groups for EC2-VPC in the Amazon EC2 User Guide for Linux Instances.

# Security considerations for mounting an Amazon EFS file system

The ability for an NFSv4.1 client to mount a file system, and assert any user and group ID when accessing it, depends on the ability to make a network connection to the NFS port of one of the file

system's mount targets. The ability to make this network connection is governed by a combination of the following:

- Network isolation provided by the mount targets' VPC File system mount targets cannot have public IP addresses associated with them. Only Amazon EC2 instances in the Amazon VPC or onpremises servers connected to the Amazon VPC via AWS Direct Connect can mount the Amazon EFS file systems. Other mechanisms for connecting to a VPC's private IP addresses from outside the VPC, such as VPN connections and VPC peering cannot currently be used to mount Amazon EFS file systems, so don't rely on them for file system access control.
- Network ACLs of the VPC subnets of the client and mount targets (for access from outside the mount target's subnets) In order to mount the file system, the client must be able to make a TCP connection to the NFS port of a mount target (and receive return traffic).
- Rules of the client's and mount targets' VPC security groups (for all access) In order for an EC2 instance to mount the file system, the following security group rules must be in effect:
  - The file system must have a mount target whose network interface has a security group with a rule than enables inbound connections on the NFS port from the instance, either by IP address (CIDR range) or security group. The source of the inbound NFS port security group rules on mount target network interfaces is a key element of file system access control. Note that inbound rules other than the NFS port, and any outbound rules, are of no use to file system mount target network interfaces.
  - The mounting instance must have a network interface with a security group rule that enables outbound connections to the NFS port on one of the file system's mount targets, either by IP address (CIDR range) or security group.

For more information, see Creating Mount Targets (p. 23).

# Read, Write, and Execute Permissions for Files and Directories

Example: Alice has permissions to read and write any files she wants in her personal directory on a file system, /alice. Alice is not allowed to read or write any files in Mark's personal directory on the same file system, /mark. Both Alice and Mark are allowed to read but not write files in the shared directory / share.

Files and directories in an EFS file system support standard Unix-style read/write/execute permissions based on the user and group ID asserted by the mounting NFSv4.1 client. For more information, see Network File System (NFS)–Level Users, Groups, and Permissions (p. 56).

#### Note

This layer of access control depends on trusting the NFSv4.1 client in its assertion of the user and group ID. There is no authentication of the identity of the NFSv4.1 client when establishing a mount connection, so any NFSv4.1 client that can make a network connection to the NFS port of a file system's mount target IP address may read and write the file system as the root user ID.

# Amazon EFS Limits

This section describes limitations when working with Amazon EFS.

#### Topics

- Amazon EFS Limits That Can Be Increased (p. 72)
- Amazon EFS Resource Limits (p. 72)
- Limits for Client EC2 Instances (p. 73)
- Limits for Amazon EFS File Systems (p. 73)
- Additional Considerations (p. 73)

## Amazon EFS Limits That Can Be Increased

Following are the limits for Amazon EFS that can be increased by contacting AWS Support.

Resource	Default Limit
Number of file systems per customer account in an AWS Region	10
Throughput limit for each file system for all connected clients	3 GB/s

You can take the following steps to request an increase for these limits. These increases are not granted immediately, so it might take a couple of days for your increase to become effective.

### To request a limit increase

- 1. Open the AWS Support Center page, sign in, if necessary, and then choose Create Case.
- 2. Under Regarding, choose Service Limit Increase.
- 3. Under **Limit Type**, choose the type of limit to increase, fill in the necessary fields in the form, and then choose your preferred method of contact.

# Amazon EFS Resource Limits

Following are the limits on Amazon EFS resources per customer account in an AWS Region.

Resource	Limit
Default limit on the number of file systems	10
Number of mount targets per file system per Availability Zone	1
Number of security groups per mount target	5
Number of tags per file system	50
Number of VPCs per file system	1

# Limits for Client EC2 Instances

The following limits for client EC2 instances apply, assuming a Linux NFSv4.1 client.

- The maximum throughput you can drive per Amazon EC2 instance is 250 MB/s.
- Up to 128 active user accounts per instance may have files open at the same time.
- Up to 32,768 files open at the same time on the instance.
- Each unique mount on the instance can acquire up to a total of 8,192 locks across a maximum of 256 unique file/process pairs. For example, a single process can acquire one or more locks on 256 separate files, or 8 processes can each acquire one or more locks on 32 files.
- Using Amazon EFS with Microsoft Windows Amazon EC2 instances is not supported.

# Limits for Amazon EFS File Systems

The following are limits specific to the Amazon EFS file systems:

- Each Amazon EFS file system has a default throughput limit of 3 GB/s for all connected clients.
- Maximum name length: 255 bytes.
- Maximum symbolic link (symlink) length: 4080 bytes.
- Maximum number of hard links to a file 175.
- Maximum size of a single file: 52,673,613,135,872 bytes (52 TiB).
- Maximum directory depth: 1000 levels deep.
- Any one particular file can have up to 87 locks across all users of the file system. You may mount a file system on one or more Amazon EC2 instances, but the maximum 87-lock limit for a file applies.
- In General Purpose mode, there is a limit of 7000 file system operations per second. This operations limit is calculated for all clients connected to a single file system.

# **Additional Considerations**

In addition, note the following:

- For a list of AWS Regions where you can create Amazon EFS file systems, see the AWS General Reference.
- Some AWS accounts created before 2012 might have access to Availability Zones in us-east-1 that do not support creating mount targets. If you are unable to create a mount target in the region, try

a different Availability Zone in that region. However, there are cost considerations for mounting a file system on an EC2 instance in an Availability Zone through a mount target created in another Availability Zone.

- You mount your file system from EC2 instances in your VPC via the mount targets you create in the VPC. You can also mount your file system on your EC2-Classic instances (which are not in the VPC), but you must first link them to your VPC via the ClassicLink. For more information about using ClassicLink, see ClassicLink in the Amazon EC2 User Guide for Linux Instances.
- An Amazon EFS file system can be mounted from on-premises datacenter servers using AWS Direct Connect. However, other VPC private connectivity mechanisms such as a VPN connection and VPC peering are not supported.

# Troubleshooting Amazon EFS

Following, you can find information on how to troubleshoot issues for Amazon Elastic File System (Amazon EFS). For optimal performance and to avoid a variety of known NFS client bugs, we recommend a Linux kernel that is version 4.0 or newer.

#### Topics

- Troubleshooting Amazon EFS: General Issues (p. 75)
- File Operation Errors (p. 80)
- Troubleshooting AMI and Kernel Issues (p. 81)

# **Troubleshooting Amazon EFS: General Issues**

Following, you can find information about general troubleshooting issues related to Amazon EFS. For information on performance, see Amazon EFS Performance (p. 62).

### Topics

- Mount Command Fails with "wrong fs type" Error Message (p. 76)
- Mount Command Fails with "incorrect mount option" Error Message (p. 76)
- File System Mount Hangs and Then Fails with Timeout Error (p. 76)
- File System Mount Using Mount Target DNS Name Fails (p. 77)
- EC2 Instance Hangs (p. 77)
- Mount Target Lifecycle State Is Stuck (p. 77)
- File System Mount on Windows Instance Fails (p. 77)
- Application Writing Large Amounts of Data Hangs (p. 78)
- Mount Does Not Respond (p. 78)
- Open and Close Operations Are Serialized (p. 78)
- Operations on Newly Mounted File System Return "bad file handle" Error (p. 79)

• Custom NFS Settings Causing Write Delays (p. 79)

## Mount Command Fails with "wrong fs type" Error Message

The mount command fails with the following error message.

```
mount: wrong fs type, bad option, bad superblock on 10.1.25.30:/,
missing codepage or helper program, or other error (for several filesystems
(e.g. nfs, cifs) you might need a /sbin/mount.<type> helper program)
In some cases useful info is found in syslog - try dmesg | tail or so.
```

#### Action to Take

Install the nfs-utils (or nfs-common on Ubuntu) package. For more information, see Installing the NFS Client (p. 40).

### Mount Command Fails with "incorrect mount option" Error Message

The mount command fails with the following error message.

```
mount.nfs4: an incorrect mount option was specified
```

Action to Take

This error message most likely means that your Linux distribution doesn't support Network File System version 4.1 (NFSv4.1). To confirm this is the case, you can run the following command:

\$ grep CONFIG\_NFS\_V4\_1 /boot/config\*

If the preceding command returns # CONFIG\_NFS\_V4\_1 is not set, NFSv4.1 is not supported on your Linux distribution. For a list of the Amazon Machine Images (AMIs) for Amazon Elastic Compute Cloud (Amazon EC2) that support NFSv4.1, see NFS Support (p. 39).

# File System Mount Hangs and Then Fails with Timeout Error

The file system mount command hangs for a minute or two, and then fails with a timeout error. The following code shows an example.

```
$ sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 mount-
target-ip:/ mnt
[2+ minute wait here]
mount.nfs4: Connection timed out
$Â
```

Action to Take

- This error can occur because either the Amazon EC2 instance or the mount target security groups are not configured properly. For more information, see Creating Security Groups (p. 28).
- Verify that the mount target IP address that you specified is valid. If you specify the wrong IP address and there is nothing else at that IP address to reject the mount, you can experience this issue.

## File System Mount Using Mount Target DNS Name Fails

A file system mount that is using a Domain Name Service (DNS) name for the mount target fails. The following code shows an example.

```
$ sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 mount-
target-DNS:/ mnt
mount.nfs4: Failed to resolve server mount-target-DNS:
Name or service not known.
$
```

#### Action to Take

Check your VPC configuration. If you are using a custom VPC, you need to make sure DNS settings are enabled. For more information, see Using DNS with Your VPC in the Amazon VPC User Guide.

To specify a mount target DNS name in the mount command, you must do the following:

- Connect your Amazon EC2 instance inside an Amazon VPC configured to use the DNS server provided by Amazon. For more information, see DHCP Options Sets in the Amazon VPC User Guide.
- Ensure that the Amazon VPC of the connecting Amazon EC2 instance has DNS host names enabled. For more information, see Updating DNS Support for Your VPC in the Amazon VPC User Guide.

### **EC2** Instance Hangs

Action to Take

An Amazon EC2 instance can hang because you deleted a file system mount target without first unmounting the file system. For more information about unmounting your Amazon EFS file system, see Unmounting File Systems (p. 46).

### Mount Target Lifecycle State Is Stuck

The mount target lifecycle state is stuck in the **creating** or **deleting** state.

Action to Take

Retry the CreateMountTarget or DeleteMountTarget call.

### File System Mount on Windows Instance Fails

Using Amazon EFS with Microsoft Windows Amazon EC2 instances is not supported.

## Application Writing Large Amounts of Data Hangs

An application that writes a large amount of data to Amazon EFS hangs and causes the instance to reboot.

Action to Take

If an application takes too long to write all of its data to Amazon EFS, Linux might reboot because it appears that the process has become unresponsive. Two kernel configuration parameters define this behavior, kernel.hung\_task\_panic and kernel.hung\_task\_timeout\_secs.

In the example following, the state of the hung process is reported by the ps command with D prior to the instance reboot, indicating that the process is waiting on I/O.

```
$ ps aux | grep large_io.py
root 33253 0.5 0.0 126652 5020 pts/3 D+ 18:22 0:00 python large_io.py
/efs/large_file
```

To prevent a reboot, increase the timeout period or disable kernel panics when a hung task is detected. The following command disables hung task kernel panics on most Linux systems.

```
$ sudo sysctl -w kernel.hung_task_panic=0
```

### Mount Does Not Respond

An Amazon EFS mount appears unresponsive. For example, commands like 1s hang.

Action to Take

This error can occur if another application is writing large amounts of data to the file system. Access to the files that are being written might be blocked until the operation is complete. In general, any commands or applications that attempt to access files that are being written to might appear to hang. For example, the ls command might hang when it gets to the file that is being written. This is because some Linux distributions alias the ls command so that it retrieves file attributes in addition to listing the directory contents.

To resolve this issue, verify that another application is writing files to the Amazon EFS mount, and that it is in the Uninterruptible sleep (D) state, as in the following example:

```
$ ps aux | grep large_io.py
root 33253 0.5 0.0 126652 5020 pts/3 D+ 18:22 0:00 python large_io.py /efs/
large_file
```

After you've verified that this is the case, you can address the issue by waiting for the other write operation to complete, or by implementing a workaround. In the example of 1s, you can use the /bin/ 1s command directly, instead of an alias, which will allow the command to proceed without hanging on the file being written. In general, if the application writing the data can force a data flush periodically, perhaps by using  $f_{\rm SYnC}(2)$ , this might help improve the responsiveness of your file system for other applications. However, this improvement might be at the expense of performance when the application writes data.

### **Open and Close Operations Are Serialized**

Open and close operations that are performed on a file system by a user on a single Amazon EC2 instance are serialized.

#### Action to Take

This issue can be resolved by using NFS protocol version 4.1, and an Amazon EC2 Amazon Machine Image (AMI) that includes a Linux kernel version 4.0 or newer. By using NFSv4.1 when mounting your file systems, you enable parallelized open and close operations on files. We recommend using **Amazon Linux AMI 2016.03.0** as the AMI for the Amazon EC2 instance that you mount your file system to.

If you can't use NFSv4.1, note that the Linux NFSv4.0 client serializes open and close requests by user ID and group IDs. This serialization happens even if multiple processes or multiple threads issue requests at the same time. The client only sends one open or close operation to an NFS server at a time, when all of the IDs match.

In addition, you can perform any of the following actions to resolve this issue:

- You can run each process from a different user ID on the same Amazon EC2 instance.
- You can leave the user IDs the same across all open requests, and modify the set of group IDs instead.
- You can run each process from a separate Amazon EC2 instance.

### Operations on Newly Mounted File System Return "bad file handle" Error

Operations performed on a newly mounted file system return a bad file handle error.

#### Action to Take

This error can happen if an Amazon EC2 instance was connected to one file system and one mount target with a specified IP address, and then that file system and mount target were deleted. If you create a new file system and mount target to connect to that Amazon EC2 instance with the same mount target IP address, this issue can occur. It can be resolved by unmounting the file system, and then remounting the file system on the Amazon EC2 instance. For more information about unmounting your Amazon EFS file system, see Unmounting File Systems (p. 46).

### Custom NFS Settings Causing Write Delays

You have custom NFS client settings, and it takes up to three seconds for an Amazon EC2 instance to see a write operation performed on a file system from another Amazon EC2 instance.

#### Action to Take

If you encounter this issue, you can resolve it in one of the following ways:

• If the NFS client on the Amazon EC2 instance that's reading data has attribute caching activated, unmount your file system, and remount it with the noac option to disable attribute caching. Attribute caching in NFSv4.1 is enabled by default.

#### Note

Disabling client-side caching can potentially reduce your application's performance.

• You can also clear your attribute cache on demand by using a programming language compatible with the NFS procedures. To do this, you can send an ACCESS procedure request immediately before a read request.

For example, using the Python programming language, you can construct the following call:

```
# Does an NFS ACCESS procedure request to clear the attribute cache, given a
   path to the file
```

import os
os.access(path, os.W\_OK)

# **File Operation Errors**

When you access Amazon EFS file systems, certain limits on the files in the file system apply. Exceeding these limits causes file operation errors. For more information on client and file based limits in Amazon EFS, see Limits for Client EC2 Instances (p. 73). Following, you can find some common file operation errors and the limits associated with each error.

### Command Fails with "Disk quota exceeded" Error

Amazon EFS doesn't currently support user disk quotas. This error can occur if any of the following limits have been exceeded:

- Up to 128 active user accounts can have files open at once for an instance.
- Up to 32,768 files can be open at once for an instance.
- Each unique mount on the instance can acquire up to a total of 8,192 locks across 256 unique fileprocess pairs. For example, a single process can acquire one or more locks on 256 separate files, or eight processes can each acquire one or more locks on 32 files.

Action to Take

If you encounter this issue, you can resolve it by identifying which of the preceding limits you are exceeding, and then making changes to meet that limit.

## Command Fails with "I/O error"

This error occurs when more than 128 active user accounts for each instance have files open at once.

Action to Take

If you encounter this issue, you can resolve it by reducing the number of active users that have files from your Amazon EFS file system open simultaneously on your instances to meet the supported limit.

## Command Fails with "File name is too long" Error

This error occurs when the size of a file name or its symbolic link (symlink) is too long. File names have the following limits:

- A name can be up to 255 bytes long.
- A symlink can be up to 4080 bytes in size.

Action to Take

If you encounter this issue, you can resolve it by reducing the size of your file name or symlink length to meet the supported limits.

## Command Fails with "Too many links" Error

This error occurs when there are too many hard links to a file. You can have up to 175 hard links in a file.

Action to Take

If you encounter this issue, you can resolve it by reducing the number of hard links to a file to meet the supported limit.

## Command Fails with "File too large" Error

This error occurs when a file is too large. A single file can be up to 52,673,613,135,872 bytes (52 TiB) in size.

Action to Take

If you encounter this issue, you can resolve it by reducing the size of a file to meet the supported limit.

## Command Fails with "Try again" Error

This error occurs when to many users or applications try to access a single file. When an application or user accesses a file, a lock is placed on the file. Any one particular file can have up to 87 locks among all users and applications of the file system. You can mount a file system on one or more Amazon EC2 instances, but the 87-lock limit for a file still applies.

Action to Take

If you encounter this issue, you can resolve it by reducing the number of applications or users accessing the file until that number meets the number of allowed locks or lower.

# **Troubleshooting AMI and Kernel Issues**

Following, you can find information about troubleshooting issues related to certain Amazon Machine Image (AMI) or kernel versions when using Amazon EFS from an Amazon EC2 instance.

Topics

- Unable to chown on Kernel Version 2.6.32 (p. 81)
- File System Keeps Performing Operations Repeatedly Due to Client Bug (p. 81)
- Deadlocked Client (p. 82)

## Unable to chown on Kernel Version 2.6.32

You're unable to perform a chown process on kernel version 2.6.32.

Action to Take

- If you're performing chown for the one-time setup step necessary to change ownership of the EFS root directory, you can run the chown command from an instance running a newer kernel, such as the newest version of Amazon Linux.
- If chown is part of your production workflow, you must update the kernel version to use chown.

## File System Keeps Performing Operations Repeatedly Due to Client Bug

A file system gets stuck performing repeated operations due to a client bug.

Action to Take

Update the client software to the latest version (currently, Linux kernel version 4.1).

### **Deadlocked Client**

A client becomes deadlocked. These two kernels contain the bug that causes this issue:

- CentOS-7 with kernel Linux 3.10.0-229.20.1.el7.x86\_64
- Ubuntu 15.10 with kernel Linux 4.2.0-18-generic

Action to Take

Do one of the following:

- Upgrade to a newer kernel version. For CentOS-7, kernel version Linux 3.10.0-327 or later contains the fix.
- Downgrade to an older kernel version.

# Amazon Elastic File System Walkthroughs

This section provides walkthroughs that you can use to explore Amazon EFS and test the end-to-end setup.

#### Topics

- Walkthrough 1: Create Amazon EFS File System and Mount It on an EC2 Instance Using the AWS CLI (p. 83)
- Walkthrough 2: Set Up an Apache Web Server and Serve Amazon EFS Files (p. 96)
- Walkthrough 3: Create Writable Per-User Subdirectories and Configure Automatic Remounting on Reboot (p. 101)
- Walkthrough 4: Back Up an EFS File System (p. 102)
- Walkthrough 5: Create and Mount a File System On-Premises with AWS Direct Connect (p. 114)

# Walkthrough 1: Create Amazon EFS File System and Mount It on an EC2 Instance Using the AWS CLI

This walkthrough uses the AWS CLI to explore the Amazon EFS API. In this walkthrough, you create an Amazon EFS file system, mount it on an EC2 instance in your VPC, and test the setup.

#### Note

This walkthrough is similar to the Getting Started exercise. In the Getting Started (p. 10) exercise, you use the console to create EC2 and Amazon EFS resources. In this walkthrough, you use the AWS CLI to do the same—primarily to familiarize yourself with the Amazon EFS API.

In this walkthrough, you create the following AWS resources in your account:

- Amazon EC2 resources:
  - Two security groups (for your EC2 instance and Amazon EFS file system).

You add rules to these security groups to authorize appropriate inbound/outbound access to allow your EC2 instance to connect to the file system via the mount target using a standard NFSv4.1 TCP port.

- An Amazon EC2 instance in your VPC.
- Amazon EFS resources:
  - A file system.
  - A mount target for your file system.

To mount your file system on an EC2 instance you need to create a mount target in your VPC. You can create one mount target in each of the Availability Zones in your VPC. For more information, see Amazon EFS: How It Works (p. 3).

Then, you test the file system on your EC2 instance. The cleanup step at the end of the walkthrough provides information for you to remove these resources.

The walkthrough creates all these resources in the US West (Oregon) Region (us-west-2). Whichever AWS Region you use, be sure to use it consistently. All of your resources—your VPC, EC2 resources, and Amazon EFS resources—must be in the same AWS Region.

### Before You Begin

- You can use the root credentials of your AWS account to sign in to the console and try the Getting Started exercise. However, AWS Identity and Access Management (IAM) recommends that you do not use the root credentials of your AWS account. Instead, create an administrator user in your account and use those credentials to manage resources in your account. For more information, see Setting Up (p. 8).
- You can use a default VPC or a custom VPC that you have created in your account. For this walkthrough, the default VPC configuration works. However, if you use a custom VPC, verify the following:
  - DNS hostnames are enabled. For more information, see Updating DNS Support for Your VPC in the Amazon VPC User Guide.
  - The Internet gateway is attached to your VPC. For more information, see Internet Gateways in the *Amazon VPC User Guide*.
  - The VPC subnets are configured to request public IP addresses for instances launched in the VPC subnets. For more information, see IP Addressing in Your VPC in the Amazon VPC User Guide.
  - The VPC route table includes a rule to send all Internet-bound traffic to the Internet gateway.
- You need to set up the AWS CLI and add the adminuser profile.

## Setting Up AWS CLI

Use the following instructions to set up the AWS CLI and user profile.

### To set up the AWS CLI

1. Download and configure the AWS CLI. For instructions, see the following topics in the AWS Command Line Interface User Guide.

Getting Set Up with the AWS Command Line Interface

Installing the AWS Command Line Interface

Configuring the AWS Command Line Interface

2. Set profiles.

You store user credentials in the AWS CLI config file. The example CLI commands in this walkthrough specify the adminuser profile. Create the adminuser profile in the config file. You can also set the administrator user profile as the default in the config file as shown.

```
[profile adminuser]
aws_access_key_id = admin user access key ID
aws_secret_access_key = admin user secret access key
region = us-west-2
[default]
aws_access_key_id = admin user access key ID
aws_secret_access_key = admin user secret access key
region = us-west-2
```

The preceding profile also sets the default AWS Region. If you don't specify a region in the CLI command, the us-west-2 region is assumed.

- 3. Verify the setup by entering the following command at the command prompt. Both of these commands don't provide credentials explicitly, so the credentials of the default profile are used.
  - Try the help command

You can also specify the user profile explicitly by adding the --profile parameter.

```
aws help
```

```
aws help \
--profile adminuser
```

### Next Step

Step 1: Create Amazon EC2 Resources (p. 85)

### Step 1: Create Amazon EC2 Resources

In this step, you do the following:

- Create two security groups.
- Add rules to the security groups to authorize additional access.
- Launch an EC2 instance. You create and mount an Amazon EFS file system on this instance in the next step.

#### Topics

- Step 1.1: Create Two Security Groups (p. 85)
- Step 1.2: Add Rules to the Security Groups to Authorize Inbound/Outbound Access (p. 87)
- Step 1.3: Launch an EC2 instance (p. 87)

### Step 1.1: Create Two Security Groups

In this section, you create security groups in your VPC for your EC2 instance and Amazon EFS mount target. Later in the walkthrough, you assign these security groups to an EC2 instance and an Amazon

EFS mount target. For information about security groups, see Security Groups for EC2-VPC in the *Amazon EC2 User Guide for Linux Instances*.

#### To create security groups

- 1. Create two security groups using the create-security-group CLI command.
  - a. Create a security group (efs-walkthrough1-ec2-sg) for your EC2 instance. You will need to provide your VPC ID.

```
$ aws ec2 create-security-group \
--region us-west-2 \
--group-name efs-walkthrough1-ec2-sg \
--description "Amazon EFS walkthrough 1, SG for EC2 instance" \
--vpc-id vpc-id-in-us-west-2 \
--profile adminuser
```

Write down the security group ID. The following is an example response:

```
{
    "GroupId": "sg-aexample"
}
```

You can find the VPC ID using the following command:

```
$ aws ec2 describe-vpcs
```

b. Create a security group (<code>efs-walkthrough1-mt-sg</code>) for your Amazon EFS mount target. You need to provide your VPC ID.

```
$ aws ec2 create-security-group \
--region us-west-2 \
--group-name efs-walkthrough1-mt-sg \
--description "Amazon EFS walkthrough 1, SG for mount target" \
--vpc-id vpc-id-in-us-west-2 \
--profile adminuser
```

Write down the security group ID. The following is an example response:

```
{
    "GroupId": "sg-aexample"
}
```

2. Verify the security groups.

```
aws ec2 describe-security-groups \
--group-ids list of security group IDs separated by space \
--profile adminuser \
--region us-west-2
```

Both should have only one outbound rule that allows all traffic to leave.

In the next section, you authorize additional access that enable the following:

- Enable you to connect to your EC2 instance.
- Enable traffic between an EC2 instance and an Amazon EFS mount target (to which you will associate these security groups later in this walkthrough).

# Step 1.2: Add Rules to the Security Groups to Authorize Inbound/Outbound Access

In this step, you add rules to the security groups to authorize inbound/outbound access.

### To add rules

1. Authorize incoming SSH connections to the security group for your EC2 instance (efswalkthrough1-ec2-sg) so you can connect to your EC2 instance using SSH from any host.

```
$ aws ec2 authorize-security-group-ingress \
--group-id id of the security group created for EC2 instance \
--protocol tcp \
--port 22 \
--cidr 0.0.0.0/0 \
--profile adminuser \
--region us-west-2
```

Verify that the security group has the inbound and outbound rule you added.

```
aws ec2 describe-security-groups \
--region us-west-2 \
--profile adminuser \
--group-id security-group-id
```

2. Authorize inbound access to the security group for the Amazon EFS mount target (efs-walkthrough1-mt-sg).

At the command prompt, run the following AWS CLI authorize-security-group-ingress command using the adminuser profile to add the inbound rule.

```
$ aws ec2 authorize-security-group-ingress \
--group-id ID of the security group created for Amazon EFS mount target \
--protocol tcp \
--port 2049 \
--source-group ID of the security group created for EC2 instance \
--profile adminuser \
--region us-west-2
```

3. Verify that both security groups now authorize inbound access.

```
aws ec2 describe-security-groups \
--group-names efs-walkthroughl-ec2-sg efs-walkthroughl-mt-sg \
--profile adminuser \
--region us-west-2
```

### Step 1.3: Launch an EC2 instance

In this step, you launch an EC2 instance.

### To launch an EC2 instance

- 1. Gather the following information that you need to provide when launching an EC2 instance:
  - a. Key pair name.
    - For introductory information, see Setting Up with Amazon EC2 in the Amazon EC2 User Guide for Linux Instances.
    - For instructions to create a .pem file, see Create a Key Pair in the Amazon EC2 User Guide for Linux Instances.
  - b. The AMI ID you want to launch.

The AWS CLI command you will use to launch an EC2 instance requires an AMI ID (that you want to deploy) as a parameter. The exercise uses the Amazon Linux HVM AMI.

Note

You can use most general purpose Linux-based AMIs. If you use another Linux API, keep in mind that you will use yum to install NFS client on the instance and you might need to add software packages as you need them.

For the Amazon Linux HVM AMI, you can find the latest IDs at Amazon Linux AMI. You choose the ID value from the Amazon Linux AMI IDs table as follows:

- Choose the **US West Oregon** region. This walkthrough assumes you are creating all resources in the US West (Oregon) Region (us-west-2).
- Choose the **EBS-backed HVM 64-bit** type (because in the CLI command you specify the t2.micro instance type, which does not support instance store).
- c. ID of the security group you created for an EC2 instance.
- d. AWS Region. This walkthrough uses the us-west-2 region.
- e. Your VPC subnet ID where you want to launch the instance. You can get list of subnets using the describe-subnets command.

```
$ aws ec2 describe-subnets \
--region us-west-2 \
--filters "Name=vpc-id,Values=vpc-id" \
--profile adminuser
```

After you choose subnet ID, write down the following values from the  ${\tt describe-subnets}$  result:

- **subnet ID** You need this value when you create a mount target. In this exercise, you create a mount target in the same subnet where you launch an EC2 instance.
- Availability Zone of the subnet You need this to construct your mount target DNS name, which you use to mount a file system on the EC2 instance.
- 2. Run the following AWS CLI run-instances command to launch an EC2 instance.

```
$ aws ec2 run-instances \
--image-id AMI ID \
--count 1 \
--instance-type t2.micro \
--associate-public-ip-address \
--key-name key-pair-name \
--security-group-ids ID of the security group created for EC2 instance \
--subnet-id VPC subnet ID \
--region us-west-2 \
```

```
--profile adminuser
```

- 3. Write down the instance ID returned by the run-instances command.
- 4. The EC2 instance you created must have a public DNS name that you use to connect to the EC2 instance and mount the file system on it. The public DNS name is of the form:

```
ec2-xx-xx-xxx.compute-1.amazonaws.com
```

Run the following CLI command and write down the public DNS name.

```
aws ec2 describe-instances \
--instance-ids <u>EC2 instance ID</u> \
--region us-west-2 \
--profile adminuser
```

If you don't find the public DNS name, check the configuration of the VPC in which you launched the EC2 instance. For more information, see Before You Begin (p. 84).

5. You can assign a name to the EC2 instance you created by adding a tag with the key Name and value set to the name you want to assign to the instance. Run the following AWS CLI create-tags command.

```
$ aws ec2 create-tags \
--resources EC2-instance-ID \
--tags Key=Name,Value=Provide-instance-name \
--region us-west-2 \
--profile adminuser
```

### Next Step

Step 2: Create Amazon EFS Resources (p. 89)

### Step 2: Create Amazon EFS Resources

In this step, you do the following:

- Create an Amazon EFS file system.
- Create a mount target in the Availability Zone where you have your EC2 instance launched.

Topics

- Step 2.1: Create Amazon EFS File System (p. 89)
- Step 2.2: Create a Mount Target (p. 91)

### Step 2.1: Create Amazon EFS File System

In this step, you create an Amazon EFS file system. Write down the FileSystemId to use later when you create mount targets for the file system in the next step.

#### To create a file system

- 1. Create a file system and add the optional Name tag.
  - a. At the command prompt, run the following AWS CLI create-file-system command.

```
$ aws efs create-file-system \
--creation-token FileSystemForWalkthroughl \
--region us-west-2 \
--profile adminuser
```

b. Verify the file system creation by calling the describe-file-systems CLI command.

```
$ aws efs describe-file-systems \
--region us-west-2 \
--profile adminuser
```

Here's an example response:

```
{
    "FileSystems": [
        {
            "SizeInBytes": {
                "Timestamp": 1418062014.0,
                "Value": 1024
            },
            "CreationToken": "FileSystemForWalkthrough1",
            "CreationTime": 1418062014.0,
            "FileSystemId": "fs-cda54064",
            "PerformanceMode" : "generalPurpose",
            "NumberOfMountTargets": 0,
            "LifeCycleState": "available",
            "OwnerId": "account-id"
        }
    ]
}
```

- c. Note the FileSystemId value. You need this value when you create a mount target for this file system in the next step.
- 2. (Optional) Add a tag to the file system you created using the create-tag CLI command.

You don't need to create a tag for your file system to complete this walkthrough. But you are exploring the Amazon EFS API, so let's test the Amazon EFS API for creating and managing tags. For more information, see CreateTags (p. 141).

a. Add a tag.

```
$ aws efs create-tags \
--file-system-id File-System-ID \
--tags Key=Name,Value=SomeExampleNameValue \
--region us-west-2 \
--profile adminuser
```

b. Retrieve a list of tags added to the file system by using the describe-tags CLI command.

```
$ aws efs describe-tags \
--file-system-id File-System-ID \
--region us-west-2 \
--profile adminuser
```

Amazon EFS returns tags list in the response body.

```
{
    "Tags": [
        {
            "Value": "SomeExampleNameValue",
            "Key": "Name"
        }
    ]
}
```

### Step 2.2: Create a Mount Target

In this step, you create a mount target for your file system in the Availability Zone where you have your EC2 instance launched.

- 1. Make sure you have the following information:
  - ID of the file system (for example, fs-example) for which you are creating the mount target.
  - VPC subnet ID where you launched the EC2 instance in Step 1.

For this walkthrough, you create the mount target in the same subnet in which you launched the EC2 instance, so you need the subnet ID (for example, subnet-example).

- ID of the security group you created for the mount target in the preceding step.
- 2. At the command prompt, run the following AWS CLI create-mount-target command.

```
$ aws efs create-mount-target \
--file-system-id file-system-id \
--subnet-id subnet-id \
--security-group ID-of-the security-group-created-for-mount-target \
--region us-west-2 \
--profile adminuser
```

You get this response:

```
{
   "MountTargetId": "fsmt-example",
   "NetworkInterfaceId": "eni-example",
   "FileSystemId": "fs-example",
   "PerformanceMode" : "generalPurpose",
   "LifeCycleState": "available",
   "SubnetId": "fs-subnet-example",
   "OwnerId": "account-id",
   "IpAddress": "xxx.xx.xxx"
}
```

3. You can also use the describe-mount-targets command to get descriptions of mount targets you created on a file system.

```
$ aws efs describe-mount-targets \
--file-system-id file-system-id \
--region us-west-2 \
--profile adminuser
```

Next Step

Step 3: Mount the Amazon EFS File System on the EC2 Instance and Test (p. 92)

# Step 3: Mount the Amazon EFS File System on the EC2 Instance and Test

In this step, you do the following:

Topics

- Step 3.1: Gather Information (p. 92)
- Step 3.2: Install the NFS Client on Your EC2 Instance (p. 92)
- Step 3.3: Mount File System on Your EC2 Instance and Test (p. 93)
- Next Step (p. 94)
- Install an NFS client on your EC2 instance.
- Mount the file system on your EC2 instance and test the setup.

### Step 3.1: Gather Information

Make sure you have the following information as you follow the steps in this section.

• Public DNS name of your EC2 instance in the following format:

ec2-xx-xxx-xx.aws-region.compute.amazonaws.com

• DNS name of your file system's mount target. You can construct this DNS name using the following generic form:

availability-zone.file-system-id.efs.aws-region.amazonaws.com

You created the mount target in the same Availability Zone where you launch your EC2 instance. So if you launched your EC2 instance in the us-west-2 AWS region and in the us-west-2a Availability Zone, then the mount target DNS name is:

us-west-2a.file-system-id.efs.us-west-2.amazonaws.com

The EC2 instance on which you mount the file system via the mount target can resolve the mount target DNS name to its IP address.

### Note

Amazon EFS does not require that your Amazon EC2 instance have either a public IP address or public DNS name. The requirements listed above are just for this walkthrough example to ensure that you'll be able to SSH into the instance from outside the VPC.

### Step 3.2: Install the NFS Client on Your EC2 Instance

You can connect to your EC2 instance from Windows or from a computer running Linux, or Mac OS X, or any other Unix variant.

### To install an NFS client

- 1. Connect to your EC2 instance:
  - To connect to your instance from a computer running Mac OS or Linux, you specify the .pem file to your ssh command with the -i option and the path to your private key.
  - To connect to your instance from a computer running Windows, you can use either MindTerm or PuTTY. If you plan to use PuTTY, you need to install it and use the following procedure to convert the .pem file to a .ppk file.

For more information, see the following topics in the Amazon EC2 User Guide for Linux Instances:

- Connecting to Your Linux Instance from Windows Using PuTTY
- Connecting to Your Linux Instance Using SSH
- 2. Execute the following commands on the EC2 instance, via the ssh session.
  - a. (Optional) Get updates and reboot.

```
$ sudo yum -y update
$ sudo reboot
```

After the reboot, reconnect to your EC2 instance.

b. Install the NFS client.

\$ sudo yum -y install nfs-utils

#### Note

If you choose the **Amazon Linux AMI 2016.03.0** Amazon Linux AMI when launching your Amazon EC2 instance, you won't need to install nfs-utils because it is already included in the AMI by default.

### Step 3.3: Mount File System on Your EC2 Instance and Test

Now you mount the file system on your EC2 instance.

1. Make a directory ("efs-mount-point").

\$ mkdir ~/efs-mount-point

2. Mount the Amazon EFS file system.

```
$ sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 mount-
target-DNS:/ ~/efs-mount-point
```

The EC2 instance can resolve the mount target DNS name to the IP address. You can optionally specify the IP address of the mount target directly.

```
$ sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 mount-
target-ip:/ ~/efs-mount-point
```

- 3. Now that you have the Amazon EFS file system mounted on your EC2 instance, you can create files.
  - a. Change the directory.

\$ cd ~/efs-mount-point

b. List the directory contents.

```
$ ls -al
```

It should be empty.

```
drwxr-xr-x 2 root root 4096 Dec 29 22:33 .
drwx----- 4 ec2-user ec2-user 4096 Dec 29 22:54 ..
```

c. The root directory of a file system, upon creation, is owned by and is writable by the root user, so you need to change permissions to add files.

\$ sudo chmod go+rw .

Now, if you try the ls -al command you see that the permissions have changed.

drwxrwxrwx 2 root root 4096 Dec 29 22:33 . drwx----- 4 ec2-user ec2-user 4096 Dec 29 22:54 ...

d. Create a text file.

```
$ touch test-file.txt
```

e. List directory content.

\$ ls -l

You now have successfully created and mounted an Amazon EFS file system on your EC2 instance in your VPC.

The file system you mounted will not persist across reboots. To automatically remount the directory you can use the fstab file. For more information, see Automatic Remounting on Reboot (p. 102). If you are using an Auto Scaling group to launch EC2 instances, you can also set scripts in a launch configuration. For an example, see Walkthrough 2: Set Up an Apache Web Server and Serve Amazon EFS Files (p. 96).

### Next Step

Step 4: Clean up (p. 94)

### Step 4: Clean up

If you no longer need the resources you created, you should remove them. You can do this with the CLI.

- Remove EC2 resources (the EC2 instance and the two security groups). Amazon EFS deletes the network interface when you delete the mount target.
- Remove Amazon EFS resources (file system, mount target).

#### To delete AWS resources created in this walkthrough

1. Terminate the EC2 instance you created for this walkthrough.

```
$ aws ec2 terminate-instances \
--instance-ids instance-id \
--profile adminuser
```

You can also delete EC2 resources using the console. For instructions, see Terminating an Instance in the Amazon EC2 User Guide for Linux Instances.

2. Delete the mount target.

You must delete the mount targets created for the file system before deleting the file system. You can get a list of mount targets by using the describe-mount-targets CLI command.

```
$ aws efs describe-mount-targets \
--file-system-id file-system-ID \
--profile adminuser \
--region aws-region
```

Then delete the mount target by using the delete-mount-target CLI command.

```
$ aws efs delete-mount-target \
--mount-target-id ID-of-mount-target-to-delete \
--profile adminuser \
--region aws-region
```

3. (Optional) Delete the two security groups you created. You don't pay for creating security groups.

You must delete the mount target's security group first, before deleting the EC2 instance's security group. The mount target's security group has a rule that references the EC2 security group. Therefore, you cannot first delete the EC2 instance's security group.

For instructions, see Deleting a Security Group in the Amazon EC2 User Guide for Linux Instances.

 Delete the file system by using the delete-file-system CLI command. You can get a list of your file systems by using the describe-file-systems CLI command. You can get the file system ID from the response.

```
aws efs describe-file-systems \
--profile adminuser \
--region aws-region
```

Delete the file system by providing the file system ID.

```
$ aws efs delete-file-system \
--file-system-id ID-of-file-system-to-delete \
--region aws-region \
--profile adminuser
```

# Walkthrough 2: Set Up an Apache Web Server and Serve Amazon EFS Files

You can have EC2 instances running the Apache web server serving files stored on your Amazon EFS file system. It can be one EC2 instance, or if your application needs, you can have multiple EC2 instances serving files from your Amazon EFS file system. The following procedures are described.

- Set up an Apache web server on an EC2 instance (p. 96).
- Set up an Apache web server on multiple EC2 instances by creating an Auto Scaling group (p. 98). You can create multiple EC2 instances using Auto Scaling, an AWS service that allows you to increase or decrease the number of EC2 instances in a group according to your application needs. When you have multiple web servers, you also need a load balancer to distribute request traffic among them.

#### Note

For both procedures, you create all resources in the US West (Oregon) Region (us-west-2).

### Single EC2 Instance Serving Files

Follow the steps to set up an Apache web server on one EC2 instance to serve files you create in your Amazon EFS file system.

- 1. Follow the steps in the Getting Started exercise so that you have a working configuration consisting of the following:
  - Amazon EFS file system
  - · EC2 instance
  - · File system mounted on the EC2 instance

For instructions, see Getting Started with Amazon Elastic File System (p. 10). As you follow the steps, write down the following:

- Public DNS name of the EC2 instance.
- Public DNS name of the mount target created in the same Availability Zone where you launched the EC2 instance.
- 2. (Optional) You may choose to unmount the file system from the mount point you created in the Getting Started exercise.

\$ sudo umount ~/efs-mount-point

In this walkthrough, you create another mount point for the file system.

- 3. On your EC2 instance, install the Apache web server and configure it as follows:
  - a. Connect to your EC2 instance and install the Apache web server.

\$ sudo yum -y install httpd

b. Start the service.

\$ sudo service httpd start

c. Create a mount point.

First note that the DocumentRoot in the /etc/httpd/conf/httpd.conf file points to / var/www/html (DocumentRoot "/var/www/html").

You will mount your Amazon EFS file system on a subdirectory under the document root.

i. Create a subdirectory efs-mount-point under /var/www/html.

\$ sudo mkdir /var/www/html/efs-mount-point

ii. Mount your Amazon EFS file system. You need to update the following command by providing your file system ID and AWS region (if you followed the Getting Started exercise to create a file system, the getting started assumes us-west-2 AWS Region).

```
$ sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2
$(curl -s http://169.254.169.254/latest/meta-data/placement/
availability-zone).file-system-ID.efs.aws-region.amazonaws.com:/
    /var/www/html/efs-mount-point
```

Here you dynamically construct DNS name of the mount target from the EC2 instance you are on. For more information, see Mounting on Amazon EC2 with a DNS Name (p. 41).

- 4. Test the setup.
  - a. Add a rule in the EC2 instance security group, which you created in the Getting Started exercise, to allow HTTP traffic on TCP port 80 from anywhere.

After you add the rule, the EC2 instance security group will have the following inbound rules.

indon indo	Cubbund Tags		
Edit	-	Post Prove (1)	•
type (1)	Protocol ()	Port Range (1)	source ()
	TCP	22	0.0.0/0
SSH			

For instructions, see Creating Security Groups Using the AWS Management Console (p. 28).b. Create a sample html file.

i. Change directory.

\$ cd /var/www/html/efs-mount-point

ii. Make a subdirectory for sampledir and change the ownership. And change directory so you can create files in the sampledir subdirectory.

```
$ sudo mkdir sampledir
$ sudo chown ec2-user sampledir
$ sudo chmod -R o+r sampledir
```

- \$ cd sampledir
- iii. Create a sample hello.html file.

```
$ echo "<html><h1>Hello from Amazon EFS</h1></html>" > hello.html
```

c. Open a browser window and enter the URL to access the file (it is the public DNS name of the EC2 instance followed by the file name). For example:

http://EC2-instance-public-DNS/efs-mount-point/sampledir/hello.html

Now you are serving web pages stored on an Amazon EFS file system.

#### Note

This setup does not configure the EC2 instance to automatically start httpd (web server) on boot, and also does not mount the file system on boot. In the next walkthrough, you create a launch configuration to set this up.

### Multiple EC2 Instances Serving Files

Follow the steps to serve the same content in your Amazon EFS file system from multiple EC2 instances for improved scalability or availability.

1. Follow the steps in the Getting Started (p. 10) exercise so that you have an Amazon EFS file system created and tested.

#### Important

For this walkthrough, you don't use the EC2 instance that you created in the Getting Started exercise. Instead, you launch new EC2 instances.

- 2. Create a load balancer in your VPC using the following steps.
  - 1. Define a load balancer

In the **Basic Configuration** section, select your VPC where you also create the EC2 instances on which you mount the file system.

In the **Select Subnets** section, you can select all of the available subnets or select a subnet that you want to use the mount target DNS name for. You can dynamically determine mount target DNS, and then the subnet does not matter. For details, see the cloud-config script in the next section.

2. Assign security groups

Create a new security group for the load balancer to allow HTTP access from port 80 from anywhere, as shown following:

- Type: HTTP
- Protocol: TCP
- Port Range: 80
- Source: Anywhere (0.0.0/0)

#### Note

When everything works, you can also update the EC2 instance security group inbound rule access to allow HTTP traffic only from the load balancer.

3. Configure a health check

Set the **Ping Path** value to /efs-mount-point/test.html. The efs-mount-point is the subdirectory where you have the file system mounted. You add test.html page in it later in this procedure.

#### Note

Don't add any EC2 instances. Later, you create an Auto Scaling Group in which you launch EC2 instance and specify this load balancer.

For instructions to create a load balancer, see Getting Started with Elastic Load Balancing in the *Elastic Load Balancing User Guide*.

- 3. Create an Auto Scaling group with two EC2 instances. First, you create a launch configuration describing the instances. Then, you create an Auto Scaling group by specifying the launch configuration. The following steps provide configuration information that you specify to create an Auto Scaling group from the Amazon EC2 console.
  - a. Choose Launch Configurations under AUTO SCALING from the left hand navigation.
  - b. Choose Create Auto Scaling group to launch the wizard.
  - c. Choose Create launch configuration.
  - d. From Quick Start, select the latest version of the Amazon Linux (HVM) AMI. This is same AMI you used in Step 1: Create Your EC2 Resources and Launch Your EC2 Instance (p. 11) of the Getting Started exercise.
  - e. In the Advanced section, do the following:
    - For IP Address Type, choose Assign a public IP address to every instance.
    - Copy/paste the following script in the User data box.

You must update the script by providing values for the *file-system-id* and *aws-region* (if you followed the Getting Started exercise, you created the file system in the us-west-2 region).

In the script, note the following:

- The script installs the NFS client and the Apache web server.
- The echo command writes the following entry in the /etc/fstab file identifying mount target DNS name and subdirectory on which to mount it. This entry ensures that the file gets mounted after each system reboot. Note that the mount target DNS name is dynamically constructed. For more information, see Mounting on Amazon EC2 with a DNS Name (p. 41).

Availability-Zone.file-system-ID.efs.aws-region.amazonaws.com:// var/www/html/efs-mount-point nfs4 defaults

- Creates efs-mount-point subdirectory and mounts the file system on it.
- Creates a test.html page so ELB health check can find the file (when creating a load balancer you specified this file as the ping point).

For more information about user data scripts, see Adding User Data in the Amazon EC2 User Guide for Linux Instances.

```
#cloud-config
package_upgrade: true
packages:
- nfs-utils
```

```
- httpd
runcmd:
- echo "$(curl -s http://169.254.169.254/latest/meta-data/placement/
availability-zone).file-system-id.efs.aws-region.amazonaws.com:/ /
var/www/html/efs-mount-point nfs4 defaults" >> /etc/fstab
- mkdir /var/www/html/efs-mount-point
- mount -a
- touch /var/www/html/efs-mount-point/test.html
- service httpd start
- chkconfig httpd on
```

f. For Assign a security group, choose Select an existing security group, and then choose the security group you created for the EC2 instance.

When configuring the Auto Scaling group details, use the following information:

- 1. For Group size, choose start with 2 instances. You will create two EC2 instances.
- 2. Select your VPC from the Network list.
- 3. Select a subnet in the same Availability Zone that you used when specifying the mount target ID in the User Data script when creating the launch configuration in the preceding step.
- 4. In the Advanced Details section
  - a. For Load Balancing, choose Receive traffic from Elastic Load Balancer(s), and then select the load balancer you created for this exercise.
  - b. For Health Check Type, choose ELB.

Follow the instructions to create an Auto Scaling group at Set Up a Scaled and Load-Balanced Application in the *Auto Scaling User Guide*. Use the information in the preceding tables where applicable.

4. Upon successful creation of the Auto Scaling group, you have two EC2 instances with nfs-utils and the Apache web server installed. On each instance, verify that you have the /var/www/ html/efs-mount-point subdirectory with your Amazon EFS file system mounted on it. For instructions to connect to an EC2 instance, see Step 3: Connect to Your Amazon EC2 Instance and Mount the Amazon EFS File System (p. 16).

#### Note

If you choose the **Amazon Linux AMI 2016.03.0** Amazon Linux AMI when launching your Amazon EC2 instance, you won't need to install nfs-utils because it is already included in the AMI by default.

- 5. Create a sample page (index.html).
  - a. Change directory.

\$ cd /var/www/html/efs-mount-point

b. Make a subdirectory for sampledir and change the ownership. And change directory so you can create files in the sampledir subdirectory. If you followed the preceding Single EC2 Instance Serving Files (p. 96), you already created the sampledir subdirectory, so you can skip this step.

```
$ sudo mkdir sampledir
$ sudo chown ec2-user sampledir
$ sudo chmod -R o+r sampledir
$ cd sampledir
```

c. Create a sample index.html file.

\$ echo "<html><h1>Hello from Amazon EFS</h1></html>" > index.html

6. Now you can test the setup. Using the load balancer's public DNS name, access the index.html page.

http://load balancer public DNS Name/efs-mount-point/sampledir/index.html

The load balancer sends a request to one of the EC2 instances running the Apache web server. Then, the web server serves the file that is stored in your Amazon EFS file system.

# Walkthrough 3: Create Writable Per-User Subdirectories and Configure Automatic Remounting on Reboot

After you create an Amazon EFS file system and mount it locally on your EC2 instance, it exposes an empty directory called the *file system root*. One common use case is to create a "writable" subdirectory under this file system root for each user you create on the EC2 instance, and mount it on the user's home directory. All files and subdirectories the user creates in their home directory are then created on the Amazon EFS file system.

In this walkthrough, you first create a user "mike" on your EC2 instance. You then mount an Amazon EFS subdirectory onto user mike's home directory. The walkthrough also explains how to configure automatic remounting of subdirectories if the system reboots.

Suppose you have an Amazon EFS file system created and mounted on a local directory on your EC2 instance. Let's call it *EFSroot*.

#### Note

You can follow the Getting Started (p. 10) exercise to create and mount an Amazon EFS file system on your EC2 instance.

In the following steps, you create a user (mike), create a subdirectory for the user (*EFSroot*/mike), make user mike the owner of the subdirectory, granting him full permissions, and finally mount the Amazon EFS subdirectory on the user's home directory (/home/mike).

- 1. Create user mike:
  - Log in to your EC2 instance. Using root privileges (in this case, using the sudo command), create user mike and assign a password.

```
$ sudo useradd -c "Mike Smith" mike
$ sudo passwd mike
```

This also creates a home directory, /home/mike, for the user.

- 2. Create a subdirectory under *EFSroot* for user mike:
  - a. Create subdirectory mike under *EFSroot*.

```
$ sudo mkdir /EFSroot/mike
```
You will need to replace *EFSroot* with your local directory name.

b. The root user and root group are the owners of the /mike subdirectory (you can verify this by using the ls -l command). To enable full permissions for user mike on this subdirectory, grant mike ownership of the directory.



3. Use the mount command to mount the *EFSroot*/mike subdirectory onto mike's home directory.

```
$ sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 mount-
target-DNS:/mike /home/mike
```

The *mount-target-DNS* address identifies the remote Amazon EFS file system root.

Now user mike's home directory is a subdirectory, writable by mike, in the Amazon EFS file system. If you unmount this mount target, the user can't access their EFS directory without remounting, which requires root permissions.

### Automatic Remounting on Reboot

You can use the file fstab to automatically remount your file system after any system reboots. For more information, see Mounting Automatically (p. 43).

### Walkthrough 4: Back Up an EFS File System

If you need to be able to recover from unintended changes or deletions in your Amazon EFS file systems, you need to implement a backup solution. This walkthrough presents a backup solution, with examples.

In this backup solution, you create an AWS Data Pipeline to copy data from your Amazon EFS file system (called the *production file system*) to another Amazon EFS file system (called the *backup file system*). This solution consists of AWS Data Pipeline templates that implement the following:

- Automated EFS backups based on a schedule that you define (for example, hourly, daily, weekly, or monthly).
- Automated rotation of the backups, where the oldest backup is replaced with the newest backup based on the number of backups that you want to retain.
- Quicker backups using rsync to only backup changes from one backup to the next.
- Efficient storage of backups using hard links. A *hard link* is a directory entry that associates a name with a file in a file system. This gives you the ability to perform a full restoration of data from any backup, while only storing what changed from backup to backup.

After you set up the backup solution, this walkthrough shows you how to access your backups to restore your data. This backup solution depends on running scripts that are hosted on GitHub, and is therefore subject to GitHub availability. If you'd prefer to eliminate this reliance and host the scripts in an Amazon S3 bucket instead, see Hosting the rysnc Scripts in an S3 Bucket (p. 112).

### Important

This solution requires using AWS Data Pipeline in the same AWS region as your file system. Because AWS Data Pipeline is not supported in US East (Ohio), this solution will not work in that region. We recommend that if you want to backup your file system, you use it in one of the other supported regions.

Topics

- EFS Backup Performance (p. 103)
- Considerations for EFS Backup (p. 103)
- Assumptions for EFS Backup (p. 104)
- Backing up an Amazon EFS File System (p. 104)
- Additional Resources (p. 109)

### EFS Backup Performance

When performing data backups and restorations, your file system performance is subject to Amazon EFS Performance (p. 62), including baseline and burst throughput capacity. Be aware that the throughput used by your backup solution counts towards your total file system throughput. The following table provides some recommendations for the Amazon EFS file system and Amazon EC2 instance sizes that work for this solution, assuming that your backup window is 15 minutes long.

EFS Size (30 MB Average File Size)	Daily Change Volume	Remaining Burst Hours	Minimum Number of Backup Agents
256 GB	less than 25 GB	6.75	1 - m3.medium
512 GB	less than 50 GB	7.75	1 - m3.large
1.0 TB	less than 75 GB	11.75	2 - m3.large*
1.5 TB	less than 125 GB	11.75	2 - m3.xlarge*
2.0 TB	less than 175 GB	11.75	3 - m3.large*
3.0 TB	less than 250 GB	11.75	4 - m3.xlarge*

\* These estimates are based on the assumption that data stored in an EFS file system that is 1 TB or larger is organized so that the backup can be spread across multiple backup nodes. The multiple-node example scripts divide the backup load across nodes based on the contents of the first-level directory of your EFS file system.

For example, if there are two backup nodes, one node backs up all of the even files and directories located in the first-level directory, while the odd node does the same for the odd files and directories. In another example, with six directories in the Amazon EFS file system and four backup nodes, the first node backs up the first and the fifth directories, the second node backs up the second and the sixth directories, and the third and fourth nodes back up the third and the fourth directories, respectively.

### Considerations for EFS Backup

Consider the following when you're deciding whether to implement this solution:

- This backup solution involves a number of AWS resources. For this solution, you need to create the following:
  - One production file system and one backup file system that contains a full copy of the production file system (plus, any incremental changes to your data over the backup rotation period).

- Amazon EC2 instances, whose lifecycles are managed by AWS Data Pipeline, that perform restorations and scheduled backups.
- One regularly scheduled AWS Data Pipeline for backing up data.
- An ad hoc AWS Data Pipeline for restoring backups.

When this solution is implemented, it results in billing to your account for these services. For more information, see the pricing pages for Amazon EFS, Amazon EC2, and AWS Data Pipeline.

• This is not an offline backup solution. To ensure a fully consistent and complete backup, pause any file writes to the file system or unmount the file system while the backup occurs. We recommend that you perform all backups during scheduled down time or off hours.

### Assumptions for EFS Backup

This walkthrough makes several assumptions and declares example values as follows:

- Before you get started, this walkthrough assumes that you already completed Getting Started (p. 10).
- After you've completed the Getting Started exercise, you have two security groups, a VPC subnet, and a file system mount target for the file system that you want to back up. For the rest of this walkthrough, you use the following example values:
  - The ID of the file system that you back up in this walkthrough is fs-12345678.
  - The security group for the file system that is associated with the mount target is called efs-mt-sg (sg-1111111a).
  - The security group that grants EC2 instances the ability to connect to the production EFS mount point is called efs-ec2-sg (sg-111111b).
  - The VPC subnet has the ID value of subnet-abcd1234.
  - The source file system mount target IP address for the file system that you want to back up is 10.0.1.32:/.
  - The example assumes that the production file system is a content management system serving media files with an average size of 30 MB.

The above assumptions and examples are reflected in the following initial setup diagram.



### Backing up an Amazon EFS File System

Follow the steps in this section to back up or restore your Amazon EFS file system.

### Topics

- Step 1: Create Your Backup EFS File System (p. 105)
- Step 2: Download the AWS Data Pipeline Template for Backups (p. 105)
- Step 3: Create a Data Pipeline for Backup (p. 106)
- Step 4: Access Your EFS Backups (p. 107)

### Step 1: Create Your Backup EFS File System

In this walkthrough, you create separate security groups, file systems, and mount points to separate your backups from your data source. In this first step, you create those resources.

- 1. First, create two new security groups. The example backup mount target security group is efsbackup-mt-sg (sg-9999999a), and the example security group for the EC2 instance that will access the mount target is efs-backup-ec2-sg (sg-9999999b). Remember to create these security groups in the same VPC as the EFS volume that you want to back up. In this example, the VPC associated with the subnet-abcd1234 subnet. For more information about creating security groups, see Creating Security Groups (p. 28).
- 2. Next, create a backup Amazon EFS file system. In this example, the file system ID is fsabcdefaa. For more information about creating file systems, see Creating an Amazon Elastic File System (p. 19).
- 3. Finally, create a mount point for the EFS backup file system and assume that it has the value of 10.0.1.75:/. For more information about creating mount targets, see Creating Mount Targets (p. 23).

After you've completed this first step, your setup should look similar to the following example diagram.



# Step 2: Download the AWS Data Pipeline Template for Backups

AWS Data Pipeline helps you reliably process and move data between different AWS compute and storage services at specified intervals. The AWS Data Pipeline console allows you to create preconfigured pipeline definitions, known as templates. You can use templates to get started with AWS Data Pipeline quickly. For this walkthrough, a template is provided to make the process of setting up your backup pipeline easier. When implemented, this template creates a data pipeline that launches a single Amazon EC2 instance on the schedule that you specify to back up data from the production file system to the backup file system. This template has a number of placeholder values. You provide the matching values for those placeholders in the **Parameters** section of the AWS Data Pipeline console. Download the AWS Data Pipeline template for backups at 1-Node-EFSBackupDataPipeline.json from GitHub.

#### Note

This template also references and runs a script to perform the backup commands. You can download the script before creating the pipeline to review what it does. To review the script, download efs-backup.sh from GitHub. This backup solution depends on running scripts that are hosted on GitHub and is subject to GitHub availability. If you'd prefer to eliminate this reliance and host the scripts in an Amazon S3 bucket instead, see Hosting the rysnc Scripts in an S3 Bucket (p. 112).

### Step 3: Create a Data Pipeline for Backup

Use the following procedure to create your data pipeline.

### To create a data pipeline for EFS backups

1. Open the AWS Data Pipeline console at https://console.aws.amazon.com/datapipeline/.

#### Important

Make sure that you're in the same AWS Region as your Amazon EFS file systems.

- 2. Choose Create new pipeline.
- 3. Add a Name and optional Description.
- 4. For Source, choose Import a definition and then choose Load local file.
- 5. In the file explorer, navigate to the template that you saved in Step 2: Download the AWS Data Pipeline Template for Backups (p. 105), and then choose **Open**.
- 6. In **Parameters**, provide the details for both your backup and production EFS file systems.

Parameters		
Production EFS mount target IP address.	10.0.1.32:/	:
Security group that can connect to the Production EFS mount point.	sg-111111b	.::
Interval for backups.	daily	•
Security group that can connect to the Backup EFS mount point.	sg-9999999b	
Number of backups to retain.	7	
Backup EFS mount target IP address.	10.0.1.75:/	:
VPC subnet for your backup EC2 instance (ideally the same subnet used for the production EFS mount point).	subnet-1234abcd	.::
Instance type for creating backups.	m3.medium	•
Name for the directory that will contain your backups.	backup-fs-12345678	:
Shell command to run.	wget https://raw.githubusercontent.com/awslabs/data-pipeline-	.::

7. Configure the options in **Schedule** to define your Amazon EFS backup schedule. The backup in the example runs once every day, and the backups are kept for a week. When a backup is seven days old, it is replaced with next oldest backup.

Schedule	
You can run your pipeline once or specify a	schedule. More
Run	once on pipeline activation
	on a schedule
Run every	1 day(s) 🔻
Starting	on pipeline activation
	2015-05-28 02:46 UTC (Current time is 02:48 UTC)
	YYYY-MM-DD HH:MM
Ending	never
	after 1 occurrence(s)
	2015-05-29 02:46 UTC (Current time is 02:48 UTC)
	YYYY-MM-DD HH:MM

### Note

We recommend that you specify a runtime that occurs during your off-peak hours.

- 8. Optionally, you can specify an Amazon S3 location for storing pipeline logs, configure a custom IAM role, or add tags to describe your pipeline.
- 9. When your pipeline is configured, choose Activate.

You've configured and activated your Amazon EFS backup data pipeline. For more information about AWS Data Pipeline, see the AWS Data Pipeline Developer Guide. At this stage, you can perform the backup now as a test, or you can wait until the backup is performed at the scheduled time.

### Step 4: Access Your EFS Backups

Your EFS backup has now been created, activated, and is running on the schedule you defined. This step outlines how you can access your Amazon EFS backups. Your backups are stored in the EFS backup file system that you created in the following format.

```
backup-efs-mount-target:/efs-backup-id/[backup interval].[0-backup
retention]>
```

Using the values from the example scenario, the backup of the file system is located in 10.1.0.75:/ fs-12345678/daily.[0-6], where daily.0 is the most recent backup and daily.6 is the oldest of the seven rotating backups.

Accessing your backups gives you the ability to restore data to your production file system. You can choose to restore an entire file system, or you can choose to restore individual files.

### Step 4.1: Restore an Entire EFS Backup

Restoring a backup copy of an Amazon EFS file system requires another AWS Data Pipeline, similar to the one you configured in Step 3: Create a Data Pipeline for Backup (p. 106). However, this restoration pipeline works in the reverse of the backup pipeline, and typically these restorations aren't scheduled to begin automatically. As with backups, restores can be done in parallel to meet your recovery time objective. Keep in mind that when you create a data pipeline, you need to schedule when you want it run. If you choose to run on activation, you start the restoration process immediately. We recommend that you only create a restoration pipeline when you need to do a restoration, or when you already have a specific window of time in mind.

Note that burst capacity is consumed by both the backup EFS and restoration EFS. For more information about performance, see Amazon EFS Performance (p. 62). The following procedure shows you how to create and implement your restoration pipeline.

### To create a data pipeline for EFS data restoration

Download the data pipeline template for restoring data from your backup EFS file system. This
template launches a single Amazon EC2 instance based on the specified size. It launches only
when you specify it to launch. Download the AWS Data Pipeline template for backups at 1-NodeEFSRestoreDataPipeline.json from GitHub.

#### Note

This template also references and runs a script to perform the restoration commands. You can download the script before creating the pipeline to review what it does. To review the script, download efs-restore.sh from GitHub.

2. Open the AWS Data Pipeline console at https://console.aws.amazon.com/datapipeline/.

#### Important

Make sure that you're in the same AWS Region as your Amazon EFS file systems and Amazon EC2.

- 3. Choose **Create new pipeline**.
- 4. Add a Name and optional Description.
- 5. For Source, choose Import a definition, and then choose Load local file.
- 6. In the file explorer, navigate to the template that you saved in Step 1: Create Your Backup EFS File System (p. 105), and then choose **Open**.
- 7. In **Parameters**, provide the details for both your backup and production EFS file systems.

Parameters	
Production EFS mount target IP address.	10.0.1.32:/
Security group that can connect to the Production EFS mount point.	sg-111111b .::
Instance type for performing the restore.	m3.large
Security group that can connect to the Backup EFS mount point.	sg-9999999b
Name for the directory that already contains your backups.	backup-fs-12345678
Backup number to restore (0 = the most recent backup).	0
Backup EFS mount target IP address.	10.0.1.75:/
Interval that you chose for the backup your going to restore.	daily
VPC subnet for your restoration EC2 instance (ideally the same subnet used for the backup EFS mount point).	subnet-1234abcd

- 8. Since you typically perform restorations only when you need them, you can schedule the restoration to run **once on pipeline activation**, or schedule a one-time restoration at a future time of your choosing, like during an off-peak window of time.
- 9. Optionally, you can specify an Amazon S3 location for storing pipeline logs, configure a custom IAM role, or add tags to describe your pipeline.
- 10. When your pipeline is configured, choose Activate.

You've configured and activated your Amazon EFS restoration data pipeline. Now, when you need to restore a backup to your production EFS file system, you just activate it from the AWS Data Pipeline console. For more information, see the AWS Data Pipeline Developer Guide.

### Step 4.2: Restore Individual Files from Your EFS Backups

You can restore files from your Amazon EFS file system backups by launching an Amazon EC2 instance to temporarily mount both the production and backup EFS file systems. The EC2 instance must be a member of both of the EFS client security groups (in this example, **efs-ec2-sg** and **efs-backup-clients-sg**). Both EFS mount targets can be mounted by this restoration instance. For example, a recovery EC2 instance could create the following mount points (the -o ro option is used to mount the Backup EFS as read-only to prevent accidentally modifying the backup when attempting to restore from a backup):

```
mount -t nfs source-efs-mount-target:/ /mnt/data
```

```
mount -t nfs -o ro backup-efs-mount-target:/fs-12345678/daily.0 /mnt/backup>
```

After you've mounted the targets, you can copy files from /mnt/backup to the appropriate location in /mnt/data in the terminal using the cp -p command. For example, an entire home directory (with its file system permissions) could be recursively copied with the following command.

sudo cp -rp /mnt/backup/users/my\_home /mnt/data/users/my\_home

And, you can restore a single file by running the following command.

```
sudo cp -p /mnt/backup/user/my_home/.profile /mnt/data/users/my_home/.profile
```

### Warning

When you are manually restoring individual data files, be careful that you don't accidentally modify the backup itself, otherwise you may corrupt it.

### Additional Resources

The backup solution presented in this walkthrough uses templates for AWS Data Pipeline. The templates used in Step 2: Download the AWS Data Pipeline Template for Backups (p. 105) and Step 4.1: Restore an Entire EFS Backup (p. 107) both use a single Amazon EC2 instance to perform their work. However, there's no real limit to the number of parallel instances that you can run for backing up or restoring your data in Amazon EFS file systems. This topic contains links to other AWS Data Pipeline templates configured for multiple EC2 instances that you can download and use for your backup solution. This section also contains instruction for how to modify the templates to include additional instances.

Topics

- Using Additional Templates (p. 109)
- Adding Additional Backup Instances (p. 110)
- Adding Additional Restoration Instances (p. 111)
- Hosting the rysnc Scripts in an S3 Bucket (p. 112)

### **Using Additional Templates**

You can download the following additional templates from GitHub:

- 2-Node-EFSBackupPipeline.json This template starts two parallel Amazon EC2 instances to backup your production Amazon EFS file system.
- 2-Node-EFSRestorePipeline.json This template starts two parallel Amazon EC2 instances to restore a backup of your production Amazon EFS file system.

### Adding Additional Backup Instances

You can add additional nodes to the backup templates used in this walkthrough. To add a node, modify the following sections of the 2-Node-EFSBackupDataPipeline.json template.

#### Important

If you're using additional nodes, you can't use spaces in file names and directories stored in the top-level directory. If you do, those files and directories won't be backed up or restored. All files and subdirectories that are at least one level below the top level will be backed up and restored as expected.

• Create an additional EC2Resource for each additional node you would like to create (in this example, a fourth EC2 instance).

```
{
  "id": "EC2Resource4",
  "terminateAfter": "70 Minutes",
  "instanceType": "#{myInstanceType}",
  "name": "EC2Resource4",
  "type": "Ec2Resource",
  "securityGroupIds" : [ "#{mySrcSecGroupID}","#{myBackupSecGroupID}" ],
  "subnetId": "#{mySubnetID}",
  "associatePublicIpAddress": "true"
},
```

- Create an additional data pipeline activity for each additional node (in this case, activity BackupPart4), make sure to configure the following sections:
  - Update the runsOn reference to point to the EC2Resource created previously (EC2Resource4 in the following example).
  - Increment the last two scriptArgument values to equal the backup part each node will be responsible for and the total number of nodes ("3" and "4" in the example below—the backup part is "3" for the fourth node in this example because our modulus logic needs to count starting with 0).

```
{
"id": "BackupPart4",
"name": "BackupPart4",
"runsOn": {
"ref": "EC2Resource4"
},
"command": "wget https://raw.githubusercontent.com/awslabs/data-pipeline-
samples/master/samples/EFSBackup/efs-backup-rsync.sh\nchmod a+x efs-backup-
rsync.sh\n./efs-backup-rsync.sh $1 $2 $3 $4 $5 $6 $7",
"scriptArgument": ["#{myEfsSource}","#{myEfsBackup}", "#{myInterval}",
"#{myRetainedBackups}","#{myEfsID}", "3", "4"],
"type": "ShellCommandActivity",
"dependsOn": {
"ref": "InitBackup"
},
"stage": "true"
},
```

• Increment the last value in all existing scriptArgument values to the number of nodes (in this example, "4").

```
{
"id": "BackupPart1",
```

```
. . .
"scriptArgument": ["#{myEfsSource}","#{myEfsBackup}", "#{myInterval}",
"#{myRetainedBackups}","#{myEfsID}", "1", "4"],
. . .
},
{
"id": "BackupPart2",
. . .
"scriptArgument": ["#{myEfsSource}","#{myEfsBackup}", "#{myInterval}",
"#{myRetainedBackups}","#{myEfsID}", "2", "4"],
. . .
},
{
"id": "BackupPart3",
. . .
"scriptArgument": ["#{myEfsSource}", "#{myEfsBackup}", "#{myInterval}",
"#{myRetainedBackups}","#{myEfsID}", "0", "4"],
. . .
},
```

• Update FinalizeBackup activity and add the new backup activity to the dependsOn list (BackupPart4 in this case).

```
{
"id": "FinalizeBackup", "name": "FinalizeBackup", "runsOn": { "ref":
"EC2Resource1" }, "command": "wget
https://raw.githubusercontent.com/awslabs/data-pipeline-samples/master/
samples/EFSBackup/efs-backup-end.sh\nchmod a+x
efs-backup-end.sh\n./efs-backup-end.sh $1 $2", "scriptArgument":
    ["#{myInterval}",
    "#{myEfsID}"], "type": "ShellCommandActivity", "dependsOn": [ { "ref":
    "BackupPart1" },
    { "ref": "BackupPart2" }, { "ref": "BackupPart3" }, { "ref":
    "BackupPart4" } ], "stage":
"true"
```

### Adding Additional Restoration Instances

You can add nodes to the restoration templates used in this walkthrough. To add a node, modify the following sections of the 2-Node-EFSRestorePipeline.json template.

• Create an additional EC2Resource for each additional node you would like to create (in this case, a third EC2 instance called EC2Resource3).

```
{
"id": "EC2Resource3",
"terminateAfter": "70 Minutes",
"instanceType": "#{myInstanceType}",
"name": "EC2Resource3",
"type": "Ec2Resource",
"securityGroupIds" : [ "#{mySrcSecGroupID}","#{myBackupSecGroupID}" ],
"subnetId": "#{mySubnetID}",
"associatePublicIpAddress": "true"
},
```

• Create an additional data pipeline activity for each additional node (in this case, Activity RestorePart3), make sure to configure the following sections:

- Update the runsOn reference to point to the EC2Resource created previously (in this example, EC2Resource3)
- Increment the last two scriptArgument values to equal the backup part each node is be
  responsible for and the total number of nodes (in this example, "2" and "3" in the example below

   the backup part is "3" for the fourth node in this example because our modulus logic needs to
   count starting with 0).

```
"id": "RestorePart3",
"name": "RestorePart3",
"runsOn": {
"ref": "EC2Resource3"
},
"command": "wget https://raw.githubusercontent.com/awslabs/data-pipeline-
samples/master/samples/EFSBackup/efs-restore-rsync.sh\nchmod a+x efs-
restore-rsync.sh\n./efs-backup-rsync.sh $1 $2 $3 $4 $5 $6 $7",
"scriptArgument": ["#{myEfsSource}","#{myEfsBackup}", "#{myInterval}",
"#{myBackup}","#{myEfsID}", "2", "3"],
"type": "ShellCommandActivity",
"dependsOn": {
"ref": "InitBackup"
},
"stage": "true"
},
```

• Increment the last value in all existing scriptArgument values to the number of nodes (in this example, "3").

```
{
"id": "RestorePart1",
...
"scriptArgument": ["#{myEfsSource}","#{myEfsBackup}", "#{myInterval}",
"#{myBackup}","#{myEfsID}", "1", "3"],
...
},
{
"id": "RestorePart2",
...
"scriptArgument": ["#{myEfsSource}","#{myEfsBackup}", "#{myInterval}",
"#{myBackup}","#{myEfsID}", "0", "3"],
...
},
```

### Hosting the rysnc Scripts in an S3 Bucket

This backup solution is dependent on running rsync scripts that are hosted in a GitHub repository on the internet. Therefore, this backup solution is subject to the GitHub repository being available. This means that if the GitHub repository removes these scripts, or if the GitHub website goes offline for some reason, the backup solution as implemented above will not function.

If you'd prefer to eliminate this GitHub dependency, you can choose to host the scripts in an Amazon S3 bucket that you own instead. This section outlines the steps necessary to host the scripts yourself.

### To host the rsync scripts in your own Amazon S3 bucket

- 1. **Sign Up for AWS** If you already have an AWS account, go ahead and skip to the next step. Otherwise, see Sign up for AWS (p. 8).
- 2. **Create an IAM User** If you already have an IAM User, go ahead and skip to the next step. Otherwise, see Create an IAM User (p. 8).
- 3. Create an Amazon S3 bucket If you already have a bucket that you want to host the rsync scripts in, go ahead and skip to the next step. Otherwise, see Create a Bucket in the Amazon Simple Storage Service Getting Started Guide.
- 4. **Download the rsync scripts and templates** Download all of all the rsync scripts and templates in the EFSBackup folder from GitHub. Make a note of the location on your computer where you downloaded these files.
- Upload the rsync scripts to your S3 bucket For instructions on how to upload objects into your S3 bucket, see Add an Object to a Bucket in the Amazon Simple Storage Service Getting Started Guide.
- 6. Change the permissions on the uploaded rsync scripts to allow **Everyone** to **Open/Download** them. For instructions on how to change the permissions on an object in your S3 bucket, see Editing Object Permissions in the *Amazon Simple Storage Service Console User Guide*.

Object: ef	s-backup.s	h				×
Bucket: Name: Link: Size: Last Modified: Owner: ETag: Expiry Date: Expiration Rule:	efs-backup.sh https://s3-us-west-2 3065 Wed Dec 02 14:30:20 None N/A	2.amazonaws.com/ GMT-800 2015		/efs-backup.sh		
▹ Details						1
• Permissions						
You can control ac	ccess to the bucket a	and its contents us	ing access policies.	Learn more.		1
Grantee: Even	yone	Open/Downloa	ad 🗖 View Permiss	sions 🔲 Edit Pern	nissions	x
Grantee:		Open/Downloa	ad 🗹 View Permiss	sions 🔽 Edit Pern	nissions	x
Add more pe	rmissions			Save	Cance	1

 Update your templates – Modify the wget statement in the shellCmd parameter to point to the Amazon S3 bucket where you put the startup script. Save the updated template, and use that template when you're following the procedure in Step 3: Create a Data Pipeline for Backup (p. 106).

#### Note

We recommend that you limit access to your Amazon S3 bucket to include the IAM account that will activate the AWS Data Pipeline for this backup solution. For more information, see Editing Bucket Permissions in the Amazon Simple Storage Service Console User Guide.

You are now hosting the rsync scripts for this backup solution, and your backups are no longer dependent on GitHub availability.

### Walkthrough 5: Create and Mount a File System On-Premises with AWS Direct Connect

This walkthrough uses the console to create and mount a file system on an on-premises server using a AWS Direct Connect connection.

In this walkthrough, it's assumed that you already have an AWS Direct Connect connection. If you don't have one, you can begin the process now and come back to this walkthrough when your connection is established. For more information, see AWS Direct Connect Product Details.

When you have an AWS Direct Connect connection, you'll create the following AWS resources in your account:

- Amazon EFS resources:
  - A file system.
  - A mount target for your file system.

To mount your file system on your on-premises servers, you need to create a mount target in your VPC. You can create one mount target in each of the Availability Zones in your VPC. For more information, see Amazon EFS: How It Works (p. 3).

Then, you'll test the file system from your on-premises server. The clean-up step at the end of the walkthrough provides information for you to remove these resources.

The walkthrough creates all these resources in the US West (Oregon) Region (us-west-2). Whichever AWS Region you use, be sure to use it consistently. All of your resources—your VPC, your mount target, and your Amazon EFS file system—must be in the same AWS Region.

#### Note

If your local application needs to know if the EFS file system is available, then your application should be able to point to a different mount point IP address if the first mount point becomes temporarily unavailable. In this scenario, we recommend that you have two on-premises servers connected to your file system through different Availability Zones (AZs) for higher availability.

### Before You Begin

You can use the root credentials of your AWS account to sign in to the console and try this exercise. However, AWS Identity and Access Management (IAM) Best practices recommend that you do not use the root credentials of your AWS account. Instead, create an administrator user in your account and use those credentials to manage resources in your account. For more information, see Setting Up (p. 8). You can use a default VPC or a custom VPC that you have created in your account. For this walkthrough, the default VPC configuration works. However, if you use a custom VPC, verify the following:

- The Internet gateway is attached to your VPC. For more information, see Internet Gateways in the *Amazon VPC User Guide*.
- The VPC route table includes a rule to send all Internet-bound traffic to the Internet gateway.

# Step 1: Create Your Amazon Elastic File System Resources

In this step, you create your Amazon EFS file system and mount targets.

### To create your Amazon EFS file system

- 1. Open the Amazon EFS console at https://console.aws.amazon.com/efs/.
- 2. Choose Create File System.
- 3. Choose your default VPC from the VPC list.
- 4. Select the check boxes for all of the Availability Zones. Make sure that they all have the default subnets, automatic IP addresses, and the default security groups chosen. These are your mount targets. For more information, see Creating Mount Targets (p. 23).

	Con	figure file syst	em access				
ep 2: Add tags ep 3: Review and create	An Am interfa	azon EFS file system ce called a mount targ VPC te mount targets	is accessed by EC2 instan et. Each mount target has	ces running ir i an IP addres	nside one of your VPCs. Inst s, which we assign automatic	ances connect to a file system via a ally or you can specify.	network
	that E	C2 instances across y Availability Zone	our VPC can access the fil Subnet	le system.	IP address	Security group	Idolity Zones
	۲		subnet-	defauit *	Automatic 🥒	sg (	default ×
	2 2		subnet	default +	Automatic /	\$9	default ×

- 5. Choose Next Step.
- 6. Name your file system, keep **general purpose** selected as your default performance mode, and choose **Next Step**.
- 7. Choose Create File System.
- 8. Choose your file system from the list and make a note of the **Security group** value. You'll need this value for the next step.

The file system you just created has mount targets, created in step 1.4. Each mount target has an associated security group. The security group acts as a virtual firewall that controls network traffic. If you didn't provide a security group when creating a mount target, Amazon EFS associates the default security group of the VPC with it. If you followed the above steps exactly, then your mount targets are using the default security group.

Next, you'll add a rule to the mount target's security group to allow inbound traffic to the NFS port (2049). You can use the AWS Management Console to add the rule to your mount target's security groups in your VPC.

#### To allow inbound traffic to the NFS port

#### Note

You don't need to add an outbound rule because the default outbound rule allows all traffic to leave (otherwise, you will need to add an outbound rule to open TCP connection on the NFS port, identifying the mount target security group as the destination).

- 1. Sign in to the AWS Management Console and open the Amazon EC2 console at https:// console.aws.amazon.com/ec2/.
- 2. Under NETWORK & SECURITY, choose Security Groups.
- 3. Choose the security group associated with your file system. You made a note of this at the end of Step 1: Create Your Amazon Elastic File System Resources (p. 115).
- 4. In the tabbed pane that appears below the list of security groups, choose the **Inbound** tab.
- 5. Choose Edit.
- 6. Choose Add Rule, and choose a rule of the following type:
  - Type NFS
  - Source Anywhere

We recommend that you only use the **Anywhere** source for testing. You can choose to create a custom source set to the IP address of the on-premises server, or use the console from the server itself, and choose **My IP**.

# Step 2: Mount the Amazon EFS File System on Your On-Premises Server

Open a terminal on your on-premises Linux server. To mount the Amazon EFS file system, you need the mount target **IP Address** for your Amazon EFS file system. If you created multiple mount targets for your file system, then you can choose any one of these. Once you have that and your terminal open, you can run the following command to mount your Amazon EFS file system.

```
$ sudo mount -t nfs4 -o
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 MOUNT_TARGET_IP_ADDRESS:/
efs
```

Change directories to the new directory that you created with the following command.

\$ cd efs

Make a subdirectory and change the ownership of that subdirectory to your EC2 instance user. Then, navigate to that new directory with the following commands.

```
$ sudo mkdir getting-started
$ sudo chown ec2-user getting-started
$ cd getting-started
```

Create a text file with the following command.

\$ touch test-file.txt

List the directory contents with the following command.

\$ ls -al

As a result, the following file is created.

```
-rw-rw-r-- 1 username username 0 Nov 15 15:32 test-file.txt
```

# Step 3: Clean Up Resources and Protect Your AWS Account

After you have finished this walkthrough, or if you don't want to explore the walkthroughs, you should follow these steps to clean up your resources and protect your AWS account.

### To clean up resources and protect your AWS account

Open the Amazon EC2 console at https://console.aws.amazon.com/ec2/.

1. Unmount the Amazon EFS file system with the following command.

\$ sudo umount efs

- 2. Open the Amazon EFS console at https://console.aws.amazon.com/efs/.
- 3. Choose the Amazon EFS file system that you want to delete from the list of file systems.
- 4. For Actions, choose Delete file system.
- 5. In the **Permanently delete file system** dialog box, type the file system ID for the Amazon EFS file system that you want to delete, and then choose **Delete File System**.
- 6. In the navigation pane, choose Security Groups.
- 7. Select the name of the security group that you added the rule to for this walkthrough.

#### Caution

Don't delete the default security group for your VPC.

- 8. For Actions, choose Edit inbound rules.
- 9. Choose the X at the end of the inbound rule you added, and choose **Save**.

# Authentication and Access Control for Amazon EFS

Access to Amazon EFS requires credentials that AWS can use to authenticate your requests. Those credentials must have permissions to access AWS resources, such an Amazon EFS file system or an Amazon Elastic Compute Cloud (Amazon EC2) instance. The following sections provide details on how you can use AWS Identity and Access Management (IAM) and Amazon EFS to help secure your resources by controlling who can access them.

- Authentication (p. 118)
- Access Control (p. 119)

### Authentication

You can access AWS as any of the following types of identities:

• AWS account root user – When you sign up for AWS, you provide an email address and password that is associated with your AWS account. These are your *root credentials* and they provide complete access to all of your AWS resources.

#### Important

For security reasons, we recommend that you use the root credentials only to create an *administrator user*, which is an *IAM user* with full permissions to your AWS account. Then, you can use this administrator user to create other IAM users and roles with limited permissions. For more information, see IAM Best Practices and Creating an Admin User and Group in the *IAM User Guide*.

IAM user – An IAM user is simply an identity within your AWS account that has specific custom
permissions (for example, permissions to create a file system in Amazon EFS). You can use an IAM
user name and password to sign in to secure AWS webpages like the AWS Management Console,
AWS Discussion Forums, or the AWS Support Center.

In addition to a user name and password, you can also generate access keys for each user. You can use these keys when you access AWS services programmatically, either through one of the several SDKs or by using the AWS Command Line Interface (CLI). The SDK and CLI tools use the access keys to cryptographically sign your request. If you don't use the AWS tools, you must sign the request yourself. Amazon EFS supports *Signature Version 4*, a protocol for authenticating inbound API requests. For more information about authenticating requests, see Signature Version 4 Signing Process in the AWS General Reference.

- IAM role An IAM role is another IAM identity you can create in your account that has specific permissions. It is similar to an *IAM user*, but it is not associated with a specific person. An IAM role enables you to obtain temporary access keys that can be used to access AWS services and resources. IAM roles with temporary credentials are useful in the following situations:
  - Federated user access Instead of creating an IAM user, you can use preexisting user identities from AWS Directory Service, your enterprise user directory, or a web identity provider. These are known as *federated users*. AWS assigns a role to a federated user when access is requested through an identity provider. For more information about federated users, see Federated Users and Roles in the *IAM User Guide*.
  - Cross-account access You can use an IAM role in your account to grant another AWS account permissions to access your account's resources. For an example, see Tutorial: Delegate Access Across AWS Accounts Using IAM Roles in the IAM User Guide.
  - AWS service access You can use an IAM role in your account to grant an AWS service permissions to access your account's resources. For example, you can create a role that allows Amazon Redshift to access an Amazon S3 bucket on your behalf and then load data stored in the bucket into an Amazon Redshift cluster. For more information, see Creating a Role to Delegate Permissions to an AWS Service in the *IAM User Guide*.
  - Applications running on Amazon EC2 Instead of storing access keys within the EC2 instance for use by applications running on the instance and making AWS API requests, you can use an IAM role to manage temporary credentials for these applications. To assign an AWS role to an EC2 instance and make it available to all of its applications, you can create an instance profile that is attached to the instance. An instance profile contains the role and enables programs running on the EC2 instance to get temporary credentials. For more information, see Using Roles for Applications on Amazon EC2 in the *IAM User Guide*.

### Access Control

You can have valid credentials to authenticate your requests, but unless you have permissions you cannot create or access Amazon Elastic File System resources. For example, you must have permissions to create an Amazon EFS file system.

The following sections describe how to manage permissions for Amazon Elastic File System. We recommend that you read the overview first.

• Overview of Managing Access Permissions to Your Amazon EFS Resources (p. 120)

• Using Identity-Based Policies (IAM Policies) for Amazon Elastic File System (p. 123)

### Overview of Managing Access Permissions to Your Amazon EFS Resources

Every AWS resource is owned by an AWS account, and permissions to create or access a resource are governed by permissions policies. An account administrator can attach permissions policies to IAM identities (that is, users, groups, and roles), and some services (such as AWS Lambda) also support attaching permissions policies to resources.

### Note

An *account administrator* (or administrator user) is a user with administrator privileges. For more information, see IAM Best Practices in the IAM User Guide.

When granting permissions, you decide who is getting the permissions, the resources they get permissions for, and the specific actions that you want to allow on those resources.

Topics

- Amazon Elastic File System Resources and Operations (p. 120)
- Understanding Resource Ownership (p. 120)
- Managing Access to Resources (p. 121)
- Specifying Policy Elements: Actions, Effects, and Principals (p. 122)
- Specifying Conditions in a Policy (p. 123)

# Amazon Elastic File System Resources and Operations

In Amazon Elastic File System, the primary resource is a *file system*. Amazon Elastic File System also supports additional resource types, the *mount target* and *tags*. However, for Amazon EFS, you can create mount targets and tags only in the context of an existing file system. Mount targets and tags are referred to as *subresource*.

These resources and subresources have unique Amazon Resource Names (ARNs) associated with them as shown in the following table.

Amazon EFS provides a set of operations to work with Amazon EFS resources. For a list of available operations, see Amazon Elastic File System Actions (p. 130).

### Understanding Resource Ownership

The AWS account owns the resources that are created in the account, regardless of who created the resources. Specifically, the resource owner is the AWS account of the principal entity (that is, the root account, an IAM user, or an IAM role) that authenticates the resource creation request. The following examples illustrate how this works:

- If you use the root account credentials of your AWS account to create a file system, your AWS account is the owner of the resource (in Amazon EFS, the resource is the file system).
- If you create an IAM user in your AWS account and grant permissions to create a file system to that user, the user can create a file system. However, your AWS account, to which the user belongs, owns the file system resource.

• If you create an IAM role in your AWS account with permissions to create a file system, anyone who can assume the role can create a file system. Your AWS account, to which the role belongs, owns the file system resource.

### Managing Access to Resources

A *permissions policy* describes who has access to what. The following section explains the available options for creating permissions policies.

#### Note

This section discusses using IAM in the context of Amazon Elastic File System. It doesn't provide detailed information about the IAM service. For complete IAM documentation, see What Is IAM? in the *IAM User Guide*. For information about IAM policy syntax and descriptions, see AWS IAM Policy Reference in the *IAM User Guide*.

Policies attached to an IAM identity are referred to as *identity-based* policies (IAM polices) and policies attached to a resource are referred to as *resource-based* policies. Amazon Elastic File System supports only identity-based policies (IAM policies).

Topics

- Identity-Based Policies (IAM Policies) (p. 121)
- Resource-Based Policies (p. 122)

### Identity-Based Policies (IAM Policies)

You can attach policies to IAM identities. For example, you can do the following:

- Attach a permissions policy to a user or a group in your account To grant a user permissions to create an Amazon EFS resource, such as a file system, you can attach a permissions policy to a user or group that the user belongs to.
- Attach a permissions policy to a role (grant cross-account permissions) You can attach an
  identity-based permissions policy to an IAM role to grant cross-account permissions. For example,
  the administrator in Account A can create a role to grant cross-account permissions to another AWS
  account (for example, Account B) or an AWS service as follows:
  - 1. Account A administrator creates an IAM role and attaches a permissions policy to the role that grants permissions on resources in Account A.
  - 2. Account A administrator attaches a trust policy to the role identifying Account B as the principal who can assume the role.
  - 3. Account B administrator can then delegate permissions to assume the role to any users in Account B. Doing this allows users in Account B to create or access resources in Account A. The principal in the trust policy can also be an AWS service principal if you want to grant an AWS service permissions to assume the role.

For more information about using IAM to delegate permissions, see Access Management in the IAM User Guide.

The following is an example policy that allows a user to perform the CreateFileSystem action for your AWS account.

```
"Version": "2012-10-17",
"Statement": [
{
"Sid" : "Stmt1EFSpermissions",
```

```
"Effect": "Allow",
      "Action": [
        "elasticfilesystem:CreateFileSystem",
        "elasticfilesystem:CreateMountTarget"
      ],
      "Resource": "arn:aws:elasticfilesystem:us-west-2:account-id:file-
system/*"
    },
    ł
     "Sid" : "Stmt2EC2permissions",
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeSubnets",
        "ec2:CreateNetworkInterface",
        "ec2:DescribeNetworkInterfaces"
      ],
      "Resource": "*"
    }
  ]
```

For more information about using identity-based policies with Amazon EFS, see Using Identity-Based Policies (IAM Policies) for Amazon Elastic File System (p. 123). For more information about users, groups, roles, and permissions, see Identities (Users, Groups, and Roles) in the *IAM User Guide*.

### **Resource-Based Policies**

Other services, such as Amazon S3, also support resource-based permissions policies. For example, you can attach a policy to an S3 bucket to manage access permissions to that bucket. Amazon Elastic File System doesn't support resource-based policies.

# Specifying Policy Elements: Actions, Effects, and Principals

For each Amazon Elastic File System resource (see Amazon Elastic File System Resources and Operations (p. 120)), the service defines a set of API operations (see Actions (p. 130)). To grant permissions for these API operations, Amazon EFS defines a set of actions that you can specify in a policy. For example, for the Amazon EFS file system resource, the following actions are defined: CreateFileSystem, DeleteFileSystem, and DescribeFileSystems. Note that, performing an API operation can require permissions for more than one action.

The following are the most basic policy elements:

- **Resource** In a policy, you use an Amazon Resource Name (ARN) to identify the resource to which the policy applies. For more information, see Amazon Elastic File System Resources and Operations (p. 120).
- Action You use action keywords to identify resource operations that you want to allow or deny. For example, depending on the specified Effect, elasticfilesystem:CreateFileSystem either allows or denies the user permissions to perform the Amazon Elastic File System CreateFileSystem operation.
- Effect You specify the effect when the user requests the specific action—this can be either allow or deny. If you don't explicitly grant access to (allow) a resource, access is implicitly denied. You can also explicitly deny access to a resource, which you might do to make sure that a user cannot access it, even if a different policy grants access.
- **Principal** In identity-based policies (IAM policies), the user that the policy is attached to is the implicit principal. For resource-based policies, you specify the user, account, service, or other entity

that you want to receive permissions (applies to resource-based policies only). Amazon EFS doesn't support resource-based policies.

To learn more about IAM policy syntax and descriptions, see AWS IAM Policy Reference in the IAM User Guide.

For a table showing all of the Amazon Elastic File System API actions, see Amazon EFS API Permissions: Actions, Resources, and Conditions Reference (p. 127).

### Specifying Conditions in a Policy

When you grant permissions, you can use the IAM policy language to specify the conditions when a policy should take effect. For example, you might want a policy to be applied only after a specific date. For more information about specifying conditions in a policy language, see Condition in the *IAM User Guide*.

To express conditions, you use predefined condition keys. There are no condition keys specific to Amazon Elastic File System. However, there are AWS-wide condition keys that you can use as appropriate. For a complete list of AWS-wide keys, see Available Keys for Conditions in the *IAM User Guide*.

#### Note

Do not use the aws:SourceIp AWS-wide condition for the CreateMountTarget, DeleteMountTarget, OF ModifyMountTargetSecurityGroup actions. Amazon EFS provisions mount targets by using its own IP address, not the IP address of the originating request.

### Using Identity-Based Policies (IAM Policies) for Amazon Elastic File System

This topic provides examples of identity-based policies that demonstrate how an account administrator can attach permissions policies to IAM identities (that is, users, groups, and roles) and thereby grant permissions to perform operations on Amazon EFS resources.

#### Important

We recommend that you first review the introductory topics that explain the basic concepts and options available for you to manage access to your Amazon Elastic File System resources. For more information, see Overview of Managing Access Permissions to Your Amazon EFS Resources (p. 120).

The sections in this topic cover the following:

- Permissions Required to Use the Amazon Elastic File System Console (p. 124)
- AWS Managed (Predefined) Policies for Amazon Elastic File System (p. 125)
- Customer Managed Policy Examples (p. 125)

The following shows an example of a permissions policy.

```
{
   "Version": "2012-10-17",
   "Statement": [
    {
        "Sid" : "AllowFileSystemPermissions",
        "Effect": "Allow",
```

```
"Action": [
        "elasticfilesystem:CreateFileSystem",
        "elasticfilesystem:CreateMountTarget"
      ],
      "Resource": "arn:aws:elasticfilesystem:us-west-2:account-id:file-
system/*"
    },
    {
     "Sid" : "AllowEC2Permissions",
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeSubnets",
        "ec2:CreateNetworkInterface",
        "ec2:DescribeNetworkInterfaces"
      ],
      "Resource": "*"
    }
  ]
}
```

The policy has two statements:

{

- The first statement grants permissions for two Amazon EFS actions (elasticfilesystem:CreateFileSystem and elasticfilesystem:CreateMountTarget) on a resource using the Amazon Resource Name (ARN) for the file system. The ARN specifies a wildcard character (\*) because you don't know the file system ID until after you create a file system.
- The second statement grants permissions for some of the Amazon EC2 actions because the elasticfilesystem:CreateMountTarget action in the first statement requires permissions for specific Amazon EC2 actions. Because these Amazon EC2 actions don't support resourcelevel permissions, the policy specifies the wildcard character (\*) as the Resource value instead of specifying a file system ARN.

The policy doesn't specify the Principal element because in an identity-based policy you don't specify the principal who gets the permission. When you attach policy to a user, the user is the implicit principal. When you attach a permissions policy to an IAM role, the principal identified in the role's trust policy gets the permissions.

For a table showing all of the Amazon Elastic File System API actions and the resources that they apply to, see Amazon EFS API Permissions: Actions, Resources, and Conditions Reference (p. 127).

### Permissions Required to Use the Amazon Elastic File System Console

The permissions reference table lists the Amazon Elastic File System API operations and shows the required permissions for each operation. For more information about Amazon EFS API operations, see Amazon EFS API Permissions: Actions, Resources, and Conditions Reference (p. 127).

To use the Amazon Elastic File System console, you need to grant permissions for additional actions as shown in the following permissions policy:

```
"Version": "2012-10-17",
"Statement": [
    {
        "Sid" : "StmtlAddtionalEC2PermissionsForConsole",
        "Effect": "Allow",
```

```
"Action": [
    "ec2:DescribeAvailabilityZones",
    "ec2:DescribeSecurityGroups",
    "ec2:DescribeVpcs",
    "ec2:DescribeVpcAttribute"
    ],
    "Resource": "*"
    }
]
```

The Amazon EFS console needs these additional permissions for the following reasons:

- Permissions for the Amazon EFS actions enable the console to display Amazon EFS resources in the account.
- The console needs permissions for the ec2 actions to query Amazon EC2 so it can display Availability Zones, VPCs, security groups, and account attributes.

# AWS Managed (Predefined) Policies for Amazon Elastic File System

AWS addresses many common use cases by providing standalone IAM policies that are created and administered by AWS. Managed policies grant necessary permissions for common use cases so you can avoid having to investigate what permissions are needed. For more information, see AWS Managed Policies in the *IAM User Guide*.

The following AWS managed policies, which you can attach to users in your account, are specific to Amazon EFS:

- AmazonElasticFileSystemReadOnlyAccess Grants read-only access to Amazon EFS resources.
- AmazonElasticFileSystemFullAccess Grants full access to Amazon EFS resources.

#### Note

You can review these permissions policies by signing in to the IAM console and searching for specific policies there.

You can also create your own custom IAM policies to allow permissions for Amazon Elastic File System API actions. You can attach these custom policies to the IAM users or groups that require those permissions.

### **Customer Managed Policy Examples**

In this section, you can find example user policies that grant permissions for various Amazon EFS actions. These policies work when you are using AWS SDKs or the AWS CLI. When you are using the console, you need to grant additional permissions specific to the console, which is discussed in Permissions Required to Use the Amazon Elastic File System Console (p. 124).

#### Note

All examples use the us-west-2 region and contain fictitious account IDs.

#### Examples

- Example 1: Allow a User to Create a Mount Target and Tags on an Existing File System (p. 126)
- Example 2: Allow a User to Perform All Amazon EFS Actions (p. 126)

# Example 1: Allow a User to Create a Mount Target and Tags on an Existing File System

The following permissions policy grants the user permissions to create mount targets and tags on a particular file system in the us-west-2 region. To create mount targets, permissions for specific Amazon EC2 actions are also required and are included in the permissions policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid" : "StmtlCreateMountTargetAndTag",
      "Effect": "Allow",
      "Action": [
        "elasticfilesystem:CreateMountTarget",
        "elasticfilesystem:DescribeMountTargets",
        "elasticfilesystem:CreateTags",
        "elasticfilesystem:DescribeTags"
      ],
      "Resource": "arn:aws:elasticfilesystem:us-west-2:123456789012:file-
system/file-system-ID"
    },
    {
      "Sid" : "Stmt2AdditionalEC2PermissionsToCreateMountTarget",
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeSubnets",
        "ec2:CreateNetworkInterface",
        "ec2:DescribeNetworkInterfaces"
      ],
      "Resource": "*"
    }
  1
}
```

### Example 2: Allow a User to Perform All Amazon EFS Actions

The following permissions policy uses a wildcard character ("elasticfilesystem:\*") to allow all Amazon EFS actions in the us-west-2 region. Because some of the Amazon EFS actions also require permissions for Amazon EC2 actions, the policy also grants permissions for all those actions.

```
{
    "Version": "2012-10-17",
    "Statement": [
    {
        "Sid" : "Stmt1PermissionForAllEFSActions",
        "Effect": "Allow",
        "Action": "elasticfilesystem:*",
        "Resource": "arn:aws:elasticfilesystem:us-west-2:123456789012:file-
system/*"
    },
    {
        "Sid" : "Stmt2RequiredEC2PermissionsForAllEFSActions",
        "Effect": "Allow",
        "Action": [
        "ec2:DescribeSubnets",
    }
}
```

```
"ec2:CreateNetworkInterface",
    "ec2:DescribeNetworkInterfaces",
    "ec2:DeleteNetworkInterfaceAttribute",
    "ec2:ModifyNetworkInterfaceAttribute",
    "ec2:DescribeNetworkInterfaceAttribute"
],
    "Resource": "*"
}
```

### Amazon EFS API Permissions: Actions, Resources, and Conditions Reference

When you are setting up Access Control (p. 119) and writing a permissions policy that you can attach to an IAM identity (identity-based policies), you can use the following list as a reference. The list includes each Amazon EFS API operation, the corresponding actions for which you can grant permissions to perform the action, and the AWS resource for which you can grant the permissions. You specify the actions in the policy's Action field, and you specify the resource value in the policy's Resource field.

You can use AWS-wide condition keys in your Amazon EFS policies to express conditions. For a complete list of AWS-wide keys, see Available Keys in the *IAM User Guide*.

#### Note

}

To specify an action, use the elasticfilesystem: prefix followed by the API operation name (for example, elasticfilesystem:CreateFileSystem).

#### Amazon EFS API and Required Permissions for Actions

```
CreateFileSystem (p. 131)
```

Action(s): elasticfilesystem:CreateFileSystem

**Resource:** arn:aws:elasticfilesystem:region:account-id:file-system/\*

#### CreateMountTarget (p. 135)

**Action(s):** elasticfilesystem:CreateMountTarget, ec2:DescribeSubnets, ec2:DescribeNetworkInterfaces, ec2:CreatenetworkInterface

**Resource:** arn:aws:elasticfilesystem:region:account-id:file-system/filesystem-id

#### CreateTags (p. 141)

Action(s): elasticfilesystem:CreateTags

Resource: arn:aws:elasticfilesystem:region:account-id:file-

system/filesystem-id

#### DeleteFileSystem (p. 143)

Action(s): elasticfilesystem: DeleteFileSystem

Resource: arn:aws:elasticfilesystem:region:account-id:file-

system/filesystem-id

#### DeleteMountTarget (p. 145)

Action(s): elasticfilesystem: DeleteMountTarget, ec2:DeleteNetworkInterface

**Resource:** arn:aws:elasticfilesystem:*region:account-id*:filesystem/*filesystem-id* 

```
DeleteTags (p. 147)
   Action(s): elasticfilesystem:DeleteTags
   Resource: arn:aws:elasticfilesystem:region:account-id:file-
   system/filesystem-id
DescribeFileSystems (p. 149)
   Action(s): elasticfilesystem: DescribeFileSystems
   Resource: arn:aws:elasticfilesystem:region:account-id:file-
   system/filesystem-id, arn:aws:elasticfilesystem:region:account-id:file-
   system/*
DescribeMountTargetSecurityGroups (p. 155)
   Action(s): elasticfilesystem: DescribeMountTargetSecurityGroups,
   ec2:DescribeNetworkInterfaceAttribute
   Resource: arn:aws:elasticfilesystem:region:account-id:file-
   system/filesystem-id
DescribeMountTargets (p. 152)
   Action(s): elasticfilesystem:DescribeMountTargets
   Resource: arn:aws:elasticfilesystem:region:account-id:file-
   system/filesystem-id
DescribeTags (p. 157)
   Action(s): elasticfilesystem:DescribeTags
   Resource: arn:aws:elasticfilesystem:region:account-id:file-
   system/filesystem-id
ModifyMountTargetSecurityGroups (p. 160)
   Action(s): elasticfilesystem: ModifyMountTargetSecurityGroups,
   ec2:ModifyNetworkInterfaceAttribute
   Resource: arn:aws:elasticfilesystem:region:account-id:file-
```

system/filesystem-id

## Amazon EFS API

The Amazon EFS API is a network protocol based on HTTP (RFC 2616). For each API call, you make an HTTP request to the region-specific Amazon EFS 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 EFS API is an RPC model, in which there is a fixed set of operations and the syntax for each operation is known to clients without any prior interaction. This section describes each API operation using an abstract RPC notation, with an operation name that does not appear on the wire. For each operation, the topic specifies the mapping to HTTP request elements.

The specific Amazon EFS 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 EFS extracts call arguments from the Request-URI path segment, query parameters, and the JSON object in the request body.

### Note

Although the operation name, such as CreateFileSystem, does not appear on the wire these names are meaningful in IAM policies. For more information, see Authentication and Access Control for Amazon EFS (p. 118). The operation name is also used to name commands in command-line tools and elements of the AWS SDK APIs. For example, there is a CLI command create-file-system that maps to the CreateFileSystem operation. It also appears in CloudTrail logs for Amazon EFS 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 region-specific and take the following form:

elasticfilesystem.aws-region.amazonaws.com

For example, the Amazon EFS API endpoint for the US West (Oregon) Region is:

elasticfilesystem.us-west-2.amazonaws.com

For a list of AWS Regions that Amazon EFS supports (where you can create and manage file systems), see Amazon Elastic File System (Amazon EFS) in the AWS General Reference.

The region-specific API endpoint defines the scope of the Amazon EFS 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 a ISO 8601 date. For example, see CreateFileSystem (p. 131).

The documentation describes API version 2015-02-01.

### **Related Topics**

The following sections provide descriptions of the API operations, how to create a signature for request authentication, and how to grant permissions for these API operations using the IAM policies.

- Authentication and Access Control for Amazon EFS (p. 118)
- Actions (p. 130)
- Data Types (p. 161)

### Actions

The following actions are supported:

- CreateFileSystem (p. 131)
- CreateMountTarget (p. 135)
- CreateTags (p. 141)
- DeleteFileSystem (p. 143)
- DeleteMountTarget (p. 145)
- DeleteTags (p. 147)
- DescribeFileSystems (p. 149)
- DescribeMountTargets (p. 152)
- DescribeMountTargetSecurityGroups (p. 155)
- DescribeTags (p. 157)
- ModifyMountTargetSecurityGroups (p. 160)

### CreateFileSystem

Creates a new, empty file system. The operation requires a creation token in the request that Amazon EFS uses to ensure idempotent creation (calling the operation with same creation token has no effect). If a file system does not currently exist that is owned by the caller's AWS account with the specified creation token, this operation does the following:

- Creates a new, empty file system. The file system will have an Amazon EFS assigned ID, and an initial lifecycle state creating.
- Returns with the description of the created file system.

Otherwise, this operation returns a FileSystemAlreadyExists error with the ID of the existing file system.

### Note

For basic use cases, you can use a randomly generated UUID for the creation token.

The idempotent operation allows you to retry a CreateFileSystem call without risk of creating an extra file system. This can happen when an initial call fails in a way that leaves it uncertain whether or not a file system was actually created. An example might be that a transport level timeout occurred or your connection was reset. As long as you use the same creation token, if the initial call had succeeded in creating a file system, the client can learn of its existence from the FileSystemAlreadyExists error.

### 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. 149) operation, which among other things returns the file system state.

This operation also takes an optional PerformanceMode parameter that you choose for your file system. We recommend generalPurpose performance mode for most file systems. File systems using the maxIO performance mode can scale to higher levels of aggregate throughput and operations per second with a tradeoff of slightly higher latencies for most file operations. The performance mode can't be changed after the file system has been created. For more information, see Amazon EFS: Performance Modes.

After the file system is fully created, Amazon EFS sets its lifecycle state to available, at which point you can create one or more mount targets for the file system in your VPC. For more information, see CreateMountTarget (p. 135). You mount your Amazon EFS file system on an EC2 instances in your VPC via the mount target. For more information, see Amazon EFS: How it Works.

This operation requires permissions for the elasticfilesystem:CreateFileSystem action.

### **Request Syntax**

```
POST /2015-02-01/file-systems HTTP/1.1
Content-type: application/json
{
    "CreationToken": "string",
    "PerformanceMode": "string"
}
```

### **URI Request Parameters**

The request does not use any URI parameters.

### Request Body

The request accepts the following data in JSON format.

#### CreationToken (p. 131)

String of up to 64 ASCII characters. Amazon EFS uses this to ensure idempotent creation. Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: Yes

#### PerformanceMode (p. 131)

The PerformanceMode of the file system. We recommend generalPurpose performance mode for most file systems. File systems using the maxIO performance mode can scale to higher levels of aggregate throughput and operations per second with a tradeoff of slightly higher latencies for most file operations. This can't be changed after the file system has been created.

Type: String

Valid Values: generalPurpose | maxIO Required: No

### Response Syntax

```
HTTP/1.1 201
Content-type: application/json
{
   "CreationTime": number,
   "CreationToken": "string",
   "FileSystemId": "string",
   "LifeCycleState": "string",
   "Name": "string",
   "NumberOfMountTargets": number,
   "OwnerId": "string",
   "PerformanceMode": "string",
   "SizeInBytes": {
      "Timestamp": number,
      "Value": number
   }
}
```

### **Response Elements**

If the action is successful, the service sends back an HTTP 201 response. The following data is returned in JSON format by the service.

```
CreationTime (p. 132)

Time that the file system was created, in seconds (since 1970-01-01T00:00:00Z).

Type: Timestamp

CreationToken (p. 132)

Opaque string specified in the request.

Type: String
```

Length Constraints: Minimum length of 1. Maximum length of 64.

### FileSystemId (p. 132)

ID of the file system, assigned by Amazon EFS. Type: String

### LifeCycleState (p. 132)

Lifecycle phase of the file system.

Type: String

Valid Values: creating | available | deleting | deleted

#### Name (p. 132)

You can add tags to a file system, including a Name tag. For more information, see CreateTags (p. 141). If the file system has a Name tag, Amazon EFS returns the value in this field.

Type: String

Length Constraints: Maximum length of 256.

#### NumberOfMountTargets (p. 132)

Current number of mount targets that the file system has. For more information, see CreateMountTarget (p. 135).

Type: Integer

Valid Range: Minimum value of 0.

### Ownerld (p. 132)

AWS account that created the file system. If the file system was created by an IAM user, the parent account to which the user belongs is the owner.

Type: String

#### PerformanceMode (p. 132)

The PerformanceMode of the file system.

Type: String

Valid Values: generalPurpose | maxIO

#### SizelnBytes (p. 132)

Latest known metered size (in bytes) of data stored in the file system, in bytes, in its Value field, and the time at which that size was determined in its Timestamp field. The Timestamp value is the integer number of seconds since 1970-01-01T00:00:00Z. Note that the value does not represent the size of a consistent snapshot of the file system, but it is eventually consistent when there are no writes to the file system. That is, the value will represent actual size only if the file system is not modified for a period longer than a couple of hours. Otherwise, the value is not the exact size the file system was at any instant in time.

Type: FileSystemSize (p. 165) object

### Errors

#### BadRequest

Returned if the request is malformed or contains an error such as an invalid parameter value or a missing required parameter.

HTTP Status Code: 400

#### FileSystemAlreadyExists

Returned if the file system you are trying to create already exists, with the creation token you provided.

HTTP Status Code: 409

### FileSystemLimitExceeded

Returned if the AWS account has already created maximum number of file systems allowed per account.

HTTP Status Code: 403

#### InternalServerError

Returned if an error occurred on the server side. HTTP Status Code: 500

### Example

### Create a file system

The following example sends a POST request to create a file system in the us-west-2 region. The request specifies myFileSystem1 as the creation token.

### Sample Request

```
POST /2015-02-01/file-systems HTTP/1.1
Host: elasticfilesystem.us-west-2.amazonaws.com
x-amz-date: 20140620T215117Z
Authorization: <...>
Content-Type: application/json
Content-Length: 42
{
  "CreationToken" : "myFileSystem1",
  "PerformanceMode" : "generalPurpose"
}
```

### Sample Response

```
HTTP/1.1 201 Created
x-amzn-RequestId: 7560489e-8bc7-4a56-a09a-757ce6f4832a
Content-Type: application/json
Content-Length: 319
{
   "ownerId": "251839141158",
   "creationToken": "myFileSystem1",
   "PerformanceMode" : "generalPurpose",
   "fileSystemId":"fs-47a2c22e",
   "CreationTime": "1403301078",
   "LifeCycleState":"creating",
   "numberOfMountTargets":0,
   "sizeInBytes":{
      "value":1024,
      "timestamp":"1403301078"
   }
}
```

### CreateMountTarget

Creates a mount target for a file system. You can then mount the file system on EC2 instances via the mount target.

You can create one mount target in each Availability Zone in your VPC. All EC2 instances in a VPC within a given Availability Zone share a single mount target for a given file system. If you have multiple subnets in an Availability Zone, you create a mount target in one of the subnets. EC2 instances do not need to be in the same subnet as the mount target in order to access their file system. For more information, see Amazon EFS: How it Works.

In the request, you also specify a file system ID for which you are creating the mount target and the file system's lifecycle state must be available. For more information, see DescribeFileSystems (p. 149).

In the request, you also provide a subnet ID, which determines the following:

- VPC in which Amazon EFS creates the mount target
- · Availability Zone in which Amazon EFS creates the mount target
- IP address range from which Amazon EFS selects the IP address of the mount target (if you don't specify an IP address in the request)

After creating the mount target, Amazon EFS returns a response that includes, a MountTargetId and an IpAddress. You use this IP address when mounting the file system in an EC2 instance. You can also use the mount target's DNS name when mounting the file system. The EC2 instance on which you mount the file system via the mount target can resolve the mount target's DNS name to its IP address. For more information, see How it Works: Implementation Overview.

Note that you can create mount targets for a file system in only one VPC, and there can be only one mount target per Availability Zone. That is, if the file system already has one or more mount targets created for it, the subnet specified in the request to add another mount target must meet the following requirements:

- Must belong to the same VPC as the subnets of the existing mount targets
- Must not be in the same Availability Zone as any of the subnets of the existing mount targets

If the request satisfies the requirements, Amazon EFS does the following:

- Creates a new mount target in the specified subnet.
- Also creates a new network interface in the subnet as follows:
  - If the request provides an IpAddress, Amazon EFS assigns that IP address to the network interface. Otherwise, Amazon EFS assigns a free address in the subnet (in the same way that the Amazon EC2 CreateNetworkInterface call does when a request does not specify a primary private IP address).
  - If the request provides SecurityGroups, this network interface is associated with those security groups. Otherwise, it belongs to the default security group for the subnet's VPC.
  - Assigns the description Mount target *fsmt-id* for file system *fs-id* where *fsmt-id* is the mount target ID, and *fs-id* is the FileSystemId.
  - Sets the requesterManaged property of the network interface to true, and the requesterId value to EFS.

Each Amazon EFS mount target has one corresponding requester-managed EC2 network interface. After the network interface is created, Amazon EFS sets the NetworkInterfaceId field in the mount target's description to the network interface ID, and the IpAddress field to its address. If network interface creation fails, the entire CreateMountTarget operation fails.

### Note

The CreateMountTarget call returns only after creating the network interface, but while the mount target state is still creating, you can check the mount target creation status by calling the DescribeMountTargets (p. 152) operation, which among other things returns the mount target state.

We recommend you create a mount target in each of the Availability Zones. There are cost considerations for using a file system in an Availability Zone through a mount target created in another Availability Zone. For more information, see Amazon EFS. In addition, by always using a mount target local to the instance's Availability Zone, you eliminate a partial failure scenario. If the Availability Zone in which your mount target is created goes down, then you won't be able to access your file system through that mount target.

This operation requires permissions for the following action on the file system:

• elasticfilesystem:CreateMountTarget

This operation also requires permissions for the following Amazon EC2 actions:

- ec2:DescribeSubnets
- ec2:DescribeNetworkInterfaces
- ec2:CreateNetworkInterface

### **Request Syntax**

```
POST /2015-02-01/mount-targets HTTP/1.1
Content-type: application/json
{
    "FileSystemId": "string",
    "IpAddress": "string",
    "SecurityGroups": [ "string" ],
    "SubnetId": "string"
}
```

### **URI Request Parameters**

The request does not use any URI parameters.

### **Request Body**

The request accepts the following data in JSON format.

### FileSystemId (p. 136)

ID of the file system for which to create the mount target. Type: String Required: Yes

### IpAddress (p. 136)

Valid IPv4 address within the address range of the specified subnet. Type: String Required: No

### SecurityGroups (p. 136)

Up to five VPC security group IDs, of the form  $\tt sg-xxxxxxx$ . These must be for the same VPC as subnet specified.

Type: array of Strings

Array Members: Maximum number of 5 items. Required: No SubnetId (p. 136) ID of the subnet to add the mount target in. Type: String

Required: Yes

### **Response Syntax**

```
HTTP/1.1 200
Content-type: application/json
{
    "FileSystemId": "string",
    "IpAddress": "string",
    "LifeCycleState": "string",
    "MountTargetId": "string",
    "NetworkInterfaceId": "string",
    "OwnerId": "string",
    "SubnetId": "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. 137)

ID of the file system for which the mount target is intended.

Type: String

### IpAddress (p. 137)

Address at which the file system may be mounted via the mount target. Type: String

### LifeCycleState (p. 137)

Lifecycle state of the mount target.

Type: String

Valid Values: creating | available | deleting | deleted

### MountTargetId (p. 137)

System-assigned mount target ID.

Type: String

### NetworkInterfaceId (p. 137)

ID of the network interface that Amazon EFS created when it created the mount target. Type: String

### **Ownerld (p. 137)**

AWS account ID that owns the resource. Type: String

### SubnetId (p. 137)

ID of the mount target's subnet. Type: String
### Errors

#### BadRequest

Returned if the request is malformed or contains an error such as an invalid parameter value or a missing required parameter.

HTTP Status Code: 400

#### FileSystemNotFound

Returned if the specified FileSystemId does not exist in the requester's AWS account. HTTP Status Code: 404

#### IncorrectFileSystemLifeCycleState

Returned if the file system's life cycle state is not "created".

HTTP Status Code: 409

#### InternalServerError

Returned if an error occurred on the server side.

HTTP Status Code: 500

#### IpAddressInUse

Returned if the request specified an IpAddress that is already in use in the subnet. HTTP Status Code: 409

#### MountTargetConflict

Returned if the mount target would violate one of the specified restrictions based on the file system's existing mount targets.

HTTP Status Code: 409

#### NetworkInterfaceLimitExceeded

The calling account has reached the ENI limit for the specific AWS region. Client should try to delete some ENIs or get its account limit raised. For more information, see Amazon VPC Limits in the Amazon Virtual Private Cloud User Guide (see the Network interfaces per VPC entry in the table).

HTTP Status Code: 409

#### NoFreeAddressesInSubnet

Returned if IpAddress was not specified in the request and there are no free IP addresses in the subnet.

HTTP Status Code: 409

#### SecurityGroupLimitExceeded

Returned if the size of SecurityGroups specified in the request is greater than five. HTTP Status Code: 400

#### SecurityGroupNotFound

Returned if one of the specified security groups does not exist in the subnet's VPC.

HTTP Status Code: 400

#### SubnetNotFound

Returned if there is no subnet with ID <code>SubnetId</code> provided in the request. HTTP Status Code: 400

#### UnsupportedAvailabilityZone

HTTP Status Code: 400

### Examples

#### Example 1: Add a mount target to a file system

The following request creates a mount target for a file system. The request specifies values for only the required FileSystemId and SubnetId parameters. The request does not provide the optional IpAddress and SecurityGroups parameters. For IpAddress, the operation uses one of the

available IP addresses in the specified subnet. And, the operation uses the default security group associated with the VPC for the SecurityGroups.

#### Sample Request

```
POST /2015-02-01/mount-targets HTTP/1.1
Host: elasticfilesystem.us-west-2.amazonaws.com
x-amz-date: 20140620T221118Z
Authorization: <...>
Content-Type: application/json
Content-Length: 160
{"SubnetId": "subnet-748c5d03", "FileSystemId": "fs-e2a6438b"}
```

#### Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: c3616af3-33fa-40ad-ae0d-d3895a2c3alf
Content-Type: application/json
Content-Length: 252
{
    "MountTargetId": "fsmt-55a4413c",
    "NetworkInterfaceId": "eni-d95852af",
    "FileSystemId": "fs-e2a6438b",
    "LifeCycleState": "available",
    "SubnetId": "subnet-748c5d03",
    "OwnerId": "231243201240",
    "IpAddress": "172.31.22.183"
}
```

#### Example 2: Add a mount target to a file system

The following request specifies all the request parameters to create a mount target.

#### Sample Request

```
POST /2015-02-01/mount-targets HTTP/1.1
Host: elasticfilesystem.us-west-2.amazonaws.com
x-amz-date: 20140620T221118Z
Authorization: <...>
Content-Type: application/json
Content-Length: 160
{
    "FileSystemId":"fs-47a2c22e",
    "SubnetId":"subnet-fd04ff94",
    "IpAddress":"10.0.2.42",
    "SecurityGroups":[
        "sg-1a2b3c4d"
    ]
}
```

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

```
Content-Type: application/json
Content-Length: 252
{
    "OwnerId":"251839141158",
    "MountTargetId":"fsmt-9a13661e",
    "FileSystemId":"fs-47a2c22e",
    "SubnetId":"subnet-fd04ff94",
    "LifeCycleState":"available",
    "IpAddress":"10.0.2.42",
    "NetworkInterfaceId":"eni-1bcb7772"
}
```

## CreateTags

Creates or overwrites tags associated with a file system. Each tag is a key-value pair. If a tag key specified in the request already exists on the file system, this operation overwrites its value with the value provided in the request. If you add the Name tag to your file system, Amazon EFS returns it in the response to the DescribeFileSystems (p. 149) operation.

This operation requires permission for the elasticfilesystem:CreateTags action.

### Request Syntax

```
POST /2015-02-01/create-tags/FileSystemId HTTP/1.1
Content-type: application/json
{
    "Tags": [
        {
         "Key": "string",
         "Value": "string"
        }
    ]
}
```

### **URI Request Parameters**

The request requires the following URI parameters.

#### FileSystemId (p. 141)

ID of the file system whose tags you want to modify (String). This operation modifies the tags only, not the file system.

### **Request Body**

The request accepts the following data in JSON format.

#### Tags (p. 141)

Array of Tag objects to add. Each Tag object is a key-value pair. Type: array of Tag (p. 167) objects Required: Yes

### **Response Syntax**

HTTP/1.1 204

### **Response Elements**

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

### Errors

#### BadRequest

Returned if the request is malformed or contains an error such as an invalid parameter value or a missing required parameter. HTTP Status Code: 400

141

#### FileSystemNotFound

Returned if the specified FileSystemId does not exist in the requester's AWS account. HTTP Status Code: 404

#### InternalServerError

Returned if an error occurred on the server side. HTTP Status Code: 500

### Example

#### Create tags on a file system

The following request creates three tags ("key1", "key2", and "key3") on the specified file system.

#### Sample Request

```
POST /2015-02-01/create-tags/fs-e2a6438b HTTP/1.1
Host: elasticfilesystem.us-west-2.amazonaws.com
x-amz-date: 20140620T221118Z
Authorization: <...>
Content-Type: application/json
Content-Length: 160
{
    "Tags": [
        {
            "Value": "value1",
            "Key": "key1"
        },
        {
            "Value": "value2",
            "Key": "key2"
        },
        {
            "Value": "value3",
            "Key": "key3"
        }
    ]
}
```

#### Sample Response

HTTP/1.1 204 no content x-amzn-RequestId: c3616af3-33fa-40ad-ae0d-d3895a2c3alf

## DeleteFileSystem

Deletes a file system, permanently severing access to its contents. Upon return, the file system no longer exists and you can't access any contents of the deleted file system.

You can't delete a file system that is in use. That is, if the file system has any mount targets, you must first delete them. For more information, see DescribeMountTargets (p. 152) and DeleteMountTarget (p. 145).

#### Note

The DeleteFileSystem call returns while the file system state is still deleting. You can check the file system deletion status by calling the DescribeFileSystems (p. 149) operation, which returns a list of file systems in your account. If you pass file system ID or creation token for the deleted file system, the DescribeFileSystems (p. 149) returns a 404 FileSystemNotFound error.

This operation requires permissions for the <code>elasticfilesystem:DeleteFileSystem</code> action.

### **Request Syntax**

DELETE /2015-02-01/file-systems/FileSystemId HTTP/1.1

### **URI Request Parameters**

The request requires the following URI parameters.

FileSystemId (p. 143)

ID of the file system you want to delete.

### Request Body

The request does not have a request body.

### Response Syntax

HTTP/1.1 204

### **Response Elements**

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

### Errors

#### BadRequest

Returned if the request is malformed or contains an error such as an invalid parameter value or a missing required parameter.

HTTP Status Code: 400

#### FileSystemInUse

Returned if a file system has mount targets.

HTTP Status Code: 409

#### FileSystemNotFound

Returned if the specified FileSystemId does not exist in the requester's AWS account. HTTP Status Code: 404

#### InternalServerError

Returned if an error occurred on the server side.

HTTP Status Code: 500

### Example

### Delete a file system

The following example sends a DELETE request to the file-systems endpoint (elasticfilesystem.us-west-2.amazonaws.com/2015-02-01/file-systems/fs-47a2c22e) to delete a file system whose ID is fs-47a2c22e.

#### Sample Request

```
DELETE /2015-02-01/file-systems/fs-47a2c22e HTTP/1.1
Host: elasticfilesystem.us-west-2.amazonaws.com
x-amz-date: 20140622T233021Z
Authorization: <...>
```

```
HTTP/1.1 204 No Content
x-amzn-RequestId: a2d125b3-7ebd-4d6a-ab3d-5548630bff33
Content-Length: 0
```

## DeleteMountTarget

Deletes the specified mount target.

This operation forcibly breaks any mounts of the file system via the mount target that is being deleted, which might disrupt instances or applications using those mounts. To avoid applications getting cut off abruptly, you might consider unmounting any mounts of the mount target, if feasible. The operation also deletes the associated network interface. Uncommitted writes may be lost, but breaking a mount target using this operation does not corrupt the file system itself. The file system you created remains. You can mount an EC2 instance in your VPC via another mount target.

This operation requires permissions for the following action on the file system:

• elasticfilesystem:DeleteMountTarget

#### Note

The DeleteMountTarget call returns while the mount target state is still deleting. You can check the mount target deletion by calling the DescribeMountTargets (p. 152) operation, which returns a list of mount target descriptions for the given file system.

The operation also requires permissions for the following Amazon EC2 action on the mount target's network interface:

ec2:DeleteNetworkInterface

### **Request Syntax**

DELETE /2015-02-01/mount-targets/MountTargetId HTTP/1.1

### **URI Request Parameters**

The request requires the following URI parameters.

#### MountTargetId (p. 145)

ID of the mount target to delete (String).

### **Request Body**

The request does not have a request body.

### **Response Syntax**

HTTP/1.1 204

### **Response Elements**

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

#### Errors

#### BadRequest

Returned if the request is malformed or contains an error such as an invalid parameter value or a missing required parameter.

HTTP Status Code: 400

#### DependencyTimeout

The service timed out trying to fulfill the request, and the client should try the call again.

HTTP Status Code: 504

#### InternalServerError

Returned if an error occurred on the server side.

HTTP Status Code: 500

#### MountTargetNotFound

Returned if there is no mount target with the specified ID found in the caller's account. HTTP Status Code: 404

### Example

#### Remove a file system's mount target

The following example sends a DELETE request to delete a specific mount target.

#### Sample Request

```
DELETE /2015-02-01/mount-targets/fsmt-9al366le HTTP/1.1
Host: elasticfilesystem.us-west-2.amazonaws.com
x-amz-date: 20140622T232908Z
Authorization: <...>
```

```
HTTP/1.1 204 No Content
x-amzn-RequestId: 76787670-2797-48ee-a34f-fce2ce122fef
```

## DeleteTags

Deletes the specified tags from a file system. If the DeleteTags request includes a tag key that does not exist, Amazon EFS ignores it and doesn't cause an error. For more information about tags and related restrictions, see Tag Restrictions in the AWS Billing and Cost Management User Guide. This operation requires permissions for the elasticfilesystem:DeleteTags action.

### **Request Syntax**

```
POST /2015-02-01/delete-tags/FileSystemId HTTP/1.1
Content-type: application/json
{
    "TagKeys": [ "string" ]
}
```

### **URI Request Parameters**

The request requires the following URI parameters.

#### FileSystemId (p. 147)

ID of the file system whose tags you want to delete (String).

### Request Body

The request accepts the following data in JSON format.

#### TagKeys (p. 147)

List of tag keys to delete. Type: array of Strings Length Constraints: Minimum length of 1. Maximum length of 128. Required: Yes

### **Response Syntax**

HTTP/1.1 204

### **Response Elements**

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

### Errors

#### BadRequest

Returned if the request is malformed or contains an error such as an invalid parameter value or a missing required parameter.

HTTP Status Code: 400

#### FileSystemNotFound

Returned if the specified FileSystemId does not exist in the requester's AWS account. HTTP Status Code: 404

#### InternalServerError

Returned if an error occurred on the server side.

HTTP Status Code: 500

### Example

### Delete tags from a file system

The following request deletes the tag key2 from the tag set associated with the file system.

#### Sample Request

```
POST /2015-02-01/delete-tags/fs-e2a6438b HTTP/1.1
Host: elasticfilesystem.us-west-2.amazonaws.com
x-amz-date: 20140620T215123Z
Authorization: <...>
Content-Type: application/json
Content-Length: 223
{
    "TagKeys":[
        "key2"
    ]
}
```

#### Sample Response

HTTP/1.1 204 No Content x-amzn-RequestId: ec08ae47-3409-49f3-9e90-64a5f981bb2b

## DescribeFileSystems

Returns the description of a specific Amazon EFS file system if either the file system CreationToken or the FileSystemId is provided. Otherwise, it returns descriptions of all file systems owned by the caller's AWS account in the AWS Region of the endpoint that you're calling.

When retrieving all file system descriptions, you can optionally specify the MaxItems parameter to limit the number of descriptions in a response. If more file system descriptions remain, Amazon EFS returns a NextMarker, an opaque token, in the response. In this case, you should send a subsequent request with the Marker request parameter set to the value of NextMarker.

To retrieve a list of your file system descriptions, this operation is used in an iterative process, where DescribeFileSystems is called first without the Marker and then the operation continues to call it with the Marker parameter set to the value of the NextMarker from the previous response until the response has no NextMarker.

The implementation may return fewer than MaxItems file system descriptions while still including a NextMarker 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 multi-call iteration is unspecified.

This operation requires permissions for the elasticfilesystem:DescribeFileSystems action.

### Request Syntax

```
GET /2015-02-01/file-
```

```
systems&CreationToken=CreationToken&FileSystemId=FileSystemId&Marker=Marker?
MaxItems=MaxItems HTTP/1.1
```

### **URI Request Parameters**

The request requires the following URI parameters.

#### CreationToken (p. 149)

(Optional) Restricts the list to the file system with this creation token (String). You specify a creation token when you create an Amazon EFS file system.

Length Constraints: Minimum length of 1. Maximum length of 64.

#### FileSystemId (p. 149)

(Optional) ID of the file system whose description you want to retrieve (String).

#### Marker (p. 149)

(Optional) Opaque pagination token returned from a previous DescribeFileSystems operation (String). If present, specifies to continue the list from where the returning call had left off.

#### MaxItems (p. 149)

(Optional) Specifies the 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 EFS returns is the minimum of the MaxItems parameter specified in the request and the service's internal maximum number of items per page.

Valid Range: Minimum value of 1.

### Request Body

The request does not have a request body.

### Response Syntax

HTTP/1.1 200

```
Content-type: application/json
{
   "FileSystems": [
     {
         "CreationTime": number,
         "CreationToken": "string",
         "FileSystemId": "string",
         "LifeCycleState": "string",
         "Name": "string",
         "NumberOfMountTargets": number,
         "OwnerId": "string",
         "PerformanceMode": "string",
         "SizeInBytes": {
            "Timestamp": number,
            "Value": number
         }
      }
   ],
   "Marker": "string",
   "NextMarker": "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. 149)

Array of file system descriptions. Type: array of FileSystemDescription (p. 163) objects

#### Marker (p. 149)

Present if provided by caller in the request (String). Type: String

#### NextMarker (p. 149)

Present if there are more file systems than returned in the response (String). You can use the NextMarker in the subsequent request to fetch the descriptions. Type: String

### Errors

#### BadRequest

Returned if the request is malformed or contains an error such as an invalid parameter value or a missing required parameter.

HTTP Status Code: 400

#### FileSystemNotFound

Returned if the specified FileSystemId does not exist in the requester's AWS account. HTTP Status Code: 404

#### InternalServerError

Returned if an error occurred on the server side. HTTP Status Code: 500

### Example

#### Retrieve list of ten file systems

The following example sends a GET request to the file-systems endpoint (elasticfilesystem.us-west-2.amazonaws.com/2015-02-01/file-systems). The request specifies a MaxItems query parameter to limit the number of file system descriptions to 10.

#### Sample Request

```
GET /2015-02-01/file-systems?MaxItems=10 HTTP/1.1
Host: elasticfilesystem.us-west-2.amazonaws.com
x-amz-date: 20140622T191208Z
Authorization: <...>
```

```
HTTP/1.1 200 OK
x-amzn-RequestId: ab5f2427-3ab3-4002-868e-30a77a88f739
Content-Type: application/json
Content-Length: 499
{
   "FileSystems":[
      {
         "OwnerId":"251839141158",
         "CreationToken": "MyFileSystem1",
         "FileSystemId":"fs-47a2c22e",
         "PerformanceMode" : "generalPurpose",
         "CreationTime": "1403301078",
         "LifeCycleState":"created",
         "Name": "my first file system",
         "NumberOfMountTargets":1,
         "SizeInBytes":{
            "Value":29313417216,
            "Timestamp":"1403301078"
         }
      }
   ]
}
```

## DescribeMountTargets

Returns the descriptions of all the current mount targets, or a specific mount target, for a file system. When requesting all of the current mount targets, the order of mount targets returned in the response is unspecified.

This operation requires permissions for the elasticfilesystem:DescribeMountTargets action, on either the file system ID that you specify in FileSystemId, or on the file system of the mount target that you specify in MountTargetId.

### Request Syntax

```
GET /2015-02-01/mount-targets&FileSystemId=FileSystemId&Marker=Marker?
MaxItems=MaxItems&MountTargetId=MountTargetId HTTP/1.1
```

### **URI Request Parameters**

The request requires the following URI parameters.

#### FileSystemId (p. 152)

(Optional) ID of the file system whose mount targets you want to list (String). It must be included in your request if MountTargetId is not included.

#### Marker (p. 152)

(Optional) Opaque pagination token returned from a previous DescribeMountTargets operation (String). If present, it specifies to continue the list from where the previous returning call left off.

#### MaxItems (p. 152)

(Optional) Maximum number of mount targets to return in the response. It must be an integer with a value greater than zero.

Valid Range: Minimum value of 1.

#### MountTargetId (p. 152)

(Optional) ID of the mount target that you want to have described (String). It must be included in your request if FileSystemId is not included.

### **Request Body**

The request does not have a request body.

### **Response Syntax**

```
HTTP/1.1 200
Content-type: application/json
{
   "Marker": "string",
   "MountTargets": [
      {
         "FileSystemId": "string",
         "IpAddress": "string",
         "LifeCycleState": "string",
         "MountTargetId": "string",
         "NetworkInterfaceId": "string",
         "OwnerId": "string",
         "SubnetId": "string"
      }
   1,
   "NextMarker": "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.

#### Marker (p. 152)

If the request included the Marker, the response returns that value in this field. Type: String

#### MountTargets (p. 152)

Returns the file system's mount targets as an array of MountTargetDescription objects. Type: array of MountTargetDescription (p. 166) objects

#### NextMarker (p. 152)

If a value is present, there are more mount targets to return. In a subsequent request, you can provide Marker in your request with this value to retrieve the next set of mount targets. Type: String

### **Errors**

#### BadRequest

Returned if the request is malformed or contains an error such as an invalid parameter value or a missing required parameter.

HTTP Status Code: 400

#### FileSystemNotFound

Returned if the specified FileSystemId does not exist in the requester's AWS account. HTTP Status Code: 404

#### InternalServerError

Returned if an error occurred on the server side.

#### HTTP Status Code: 500

#### MountTargetNotFound

Returned if there is no mount target with the specified ID found in the caller's account. HTTP Status Code: 404

### Example

#### Retrieve descriptions mount targets created for a file system

The following request retrieves descriptions of mount targets created for the specified file system.

#### Sample Request

```
GET /2015-02-01/mount-targets?FileSystemId=fs-47a2c22e HTTP/1.1
Host: elasticfilesystem.us-west-2.amazonaws.com
x-amz-date: 20140622T191252Z
Authorization: <...>
```

```
HTTP/1.1 200 OK
x-amzn-RequestId: ab5f2427-3ab3-4002-868e-30a77a88f739
```

```
Content-Type: application/json
Content-Length: 357
{
    "MountTargets":[
        {
            "OwnerId":"251839141158",
            "MountTargetId":"fsmt-9a13661e",
            "FileSystemId":"fs-47a2c22e",
            "SubnetId":"subnet-fd04ff94",
            "LifeCycleState":"added",
            "IpAddress":"10.0.2.42",
            "NetworkInterfaceId":"eni-1bcb7772"
        }
    ]
}
```

## DescribeMountTargetSecurityGroups

Returns the security groups currently in effect for a mount target. This operation requires that the network interface of the mount target has been created and the lifecycle state of the mount target is not deleted.

This operation requires permissions for the following actions:

- elasticfilesystem:DescribeMountTargetSecurityGroups action on the mount target's file system.
- ec2:DescribeNetworkInterfaceAttribute action on the mount target's network interface.

### **Request Syntax**

GET /2015-02-01/mount-targets/*MountTargetId*/security-groups HTTP/1.1

### **URI Request Parameters**

The request requires the following URI parameters.

#### MountTargetId (p. 155)

ID of the mount target whose security groups you want to retrieve.

### **Request Body**

The request does not have a request body.

### Response Syntax

```
HTTP/1.1 200
Content-type: application/json
{
    "SecurityGroups": [ "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.

#### SecurityGroups (p. 155)

Array of security groups. Type: array of Strings Array Members: Maximum number of 5 items.

### Errors

#### BadRequest

Returned if the request is malformed or contains an error such as an invalid parameter value or a missing required parameter. HTTP Status Code: 400

#### IncorrectMountTargetState

Returned if the mount target is not in the correct state for the operation.

HTTP Status Code: 409

#### InternalServerError

Returned if an error occurred on the server side. HTTP Status Code: 500

#### MountTargetNotFound

Returned if there is no mount target with the specified ID found in the caller's account. HTTP Status Code: 404

### Example

#### Retrieve security groups in effect for a file system

The following example retrieves the security groups that are in effect for the network interface associated with a mount target.

#### Sample Request

```
GET /2015-02-01/mount-targets/fsmt-9a13661e/security-groups HTTP/1.1
Host: elasticfilesystem.us-west-2.amazonaws.com
x-amz-date: 20140620T223513Z
Authorization: <...>
```

#### Sample Response

}

```
HTTP/1.1 200 OK
x-amzn-RequestId: 088fb0b4-0c1d-4af7-9de1-933207fbdb46
Content-Length: 57
{
"SecurityGroups" : [
"sg-188d9f74"
]
```

## DescribeTags

Returns the tags associated with a file system. The order of tags returned in the response of one DescribeTags call and the order of tags returned across the responses of a multi-call iteration (when using pagination) is unspecified.

This operation requires permissions for the elasticfilesystem:DescribeTags action.

### **Request Syntax**

GET /2015-02-01/tags/FileSystemId/&Marker=Marker?MaxItems=MaxItems HTTP/1.1

### **URI Request Parameters**

The request requires the following URI parameters.

### FileSystemId (p. 157)

ID of the file system whose tag set you want to retrieve.

#### Marker (p. 157)

(Optional) Opaque pagination token returned from a previous DescribeTags operation (String). If present, it specifies to continue the list from where the previous call left off.

#### MaxItems (p. 157)

(Optional) Maximum number of file system tags to return in the response. It must be an integer with a value greater than zero.

Valid Range: Minimum value of 1.

### Request Body

The request does not have a request body.

### Response Syntax

### **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.

#### Marker (p. 157)

If the request included a Marker, the response returns that value in this field. Type: String

#### NextMarker (p. 157)

If a value is present, there are more tags to return. In a subsequent request, you can provide the value of NextMarker as the value of the Marker parameter in your next request to retrieve the next set of tags.

Type: String

#### Tags (p. 157)

Returns tags associated with the file system as an array of Tag objects.

Type: array of Tag (p. 167) objects

### **Errors**

#### BadRequest

Returned if the request is malformed or contains an error such as an invalid parameter value or a missing required parameter.

HTTP Status Code: 400

#### FileSystemNotFound

Returned if the specified FileSystemId does not exist in the requester's AWS account. HTTP Status Code: 404

#### InternalServerError

Returned if an error occurred on the server side. HTTP Status Code: 500

### Example

#### Retrieve tags associated with a file system

The following request retrieves tags (key-value pairs) associated with the specified file system.

#### Sample Request

```
GET /2015-02-01/tags/fs-e2a6438b/ HTTP/1.1
Host: elasticfilesystem.us-west-2.amazonaws.com
x-amz-date: 20140620T215404Z
Authorization: <...>
```

```
HTTP/1.1 200 OK
x-amzn-RequestId: f264e454-7859-4f15-8169-1c0d5b0b04f5
Content-Type: application/json
Content-Length: 288
{
    "Tags":[
        {
         "Key":"Name",
         "Value":"my first file system"
        },
        {
         "Key":"Fleet",
         "Value":"Development"
        },
        {
        }
    }
}
```

```
"Key":"Developer",
"Value":"Alice"
}
]
}
```

## ModifyMountTargetSecurityGroups

Modifies the set of security groups in effect for a mount target.

When you create a mount target, Amazon EFS also creates a new network interface. For more information, see CreateMountTarget (p. 135). This operation replaces the security groups in effect for the network interface associated with a mount target, with the SecurityGroups provided in the request. This operation requires that the network interface of the mount target has been created and the lifecycle state of the mount target is not deleted.

The operation requires permissions for the following actions:

- elasticfilesystem:ModifyMountTargetSecurityGroups action on the mount target's file system.
- ec2:ModifyNetworkInterfaceAttribute action on the mount target's network interface.

### **Request Syntax**

```
PUT /2015-02-01/mount-targets/MountTargetId/security-groups HTTP/1.1
Content-type: application/json
{
    "SecurityGroups": [ "string" ]
}
```

### **URI Request Parameters**

The request requires the following URI parameters.

#### MountTargetId (p. 160)

ID of the mount target whose security groups you want to modify.

### **Request Body**

The request accepts the following data in JSON format.

#### SecurityGroups (p. 160)

Array of up to five VPC security group IDs. Type: array of Strings Array Members: Maximum number of 5 items. Required: No

### **Response Syntax**

HTTP/1.1 204

### **Response Elements**

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

### Errors

#### BadRequest

Returned if the request is malformed or contains an error such as an invalid parameter value or a missing required parameter.

HTTP Status Code: 400
-----------------------

#### IncorrectMountTargetState

Returned if the mount target is not in the correct state for the operation.

HTTP Status Code: 409

#### InternalServerError

Returned if an error occurred on the server side.

HTTP Status Code: 500

#### MountTargetNotFound

Returned if there is no mount target with the specified ID found in the caller's account. HTTP Status Code: 404

#### SecurityGroupLimitExceeded

Returned if the size of SecurityGroups specified in the request is greater than five. HTTP Status Code: 400

#### SecurityGroupNotFound

Returned if one of the specified security groups does not exist in the subnet's VPC. HTTP Status Code: 400

### Example

#### Replace a mount target's security groups

The following example replaces security groups in effect for the network interface associated with a mount target.

#### Sample Request

```
PUT /2015-02-01/mount-targets/fsmt-9al3661e/security-groups HTTP/1.1
Host: elasticfilesystem.us-west-2.amazonaws.com
x-amz-date: 20140620T223446Z
Authorization: <...>
Content-Type: application/json
Content-Length: 57
{
    "SecurityGroups" : [
    "sg-188d9f74"
]
}
```

#### Sample Response

```
HTTP/1.1 204 No Content
x-amzn-RequestId: 088fb0b4-0c1d-4af7-9de1-933207fbdb46
```

## **Data Types**

The following data types are supported:

- FileSystemDescription (p. 163)
- FileSystemSize (p. 165)
- MountTargetDescription (p. 166)
- Tag (p. 167)

## FileSystemDescription

Description of the file system.

### Contents

#### CreationTime

Time that the file system was created, in seconds (since 1970-01-01T00:00:00Z). Type: Timestamp

Required: Yes

#### CreationToken

Opaque string specified in the request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: Yes

#### FileSystemId

ID of the file system, assigned by Amazon EFS. Type: String Required: Yes

#### LifeCycleState

Lifecycle phase of the file system.

Type: String

Valid Values: creating | available | deleting | deleted

Required: Yes

#### Name

You can add tags to a file system, including a Name tag. For more information, see CreateTags (p. 141). If the file system has a Name tag, Amazon EFS returns the value in this field.

Type: String

Length Constraints: Maximum length of 256.

Required: No

#### NumberOfMountTargets

Current number of mount targets that the file system has. For more information, see CreateMountTarget (p. 135).

Type: Integer

Valid Range: Minimum value of 0.

Required: Yes

#### Ownerld

AWS account that created the file system. If the file system was created by an IAM user, the parent account to which the user belongs is the owner.

Type: String

Required: Yes

#### PerformanceMode

The PerformanceMode of the file system.

Type: String

Valid Values: generalPurpose | maxIO

Required: Yes

#### SizeInBytes

Latest known metered size (in bytes) of data stored in the file system, in bytes, in its Value field, and the time at which that size was determined in its Timestamp field. The Timestamp value is the integer number of seconds since 1970-01-01T00:00:00Z. Note that the value does not represent the size of a consistent snapshot of the file system, but it is eventually consistent when

there are no writes to the file system. That is, the value will represent actual size only if the file system is not modified for a period longer than a couple of hours. Otherwise, the value is not the exact size the file system was at any instant in time.

Type: FileSystemSize (p. 165) object Required: Yes

## FileSystemSize

Latest known metered size (in bytes) of data stored in the file system, in its Value field, and the time at which that size was determined in its Timestamp field. Note that the value does not represent the size of a consistent snapshot of the file system, but it is eventually consistent when there are no writes to the file system. That is, the value will represent the actual size only if the file system is not modified for a period longer than a couple of hours. Otherwise, the value is not necessarily the exact size the file system was at any instant in time.

### Contents

#### Timestamp

Time at which the size of data, returned in the Value field, was determined. The value is the integer number of seconds since 1970-01-01T00:00:00Z.

Type: Timestamp

Required: No

#### Value

Latest known metered size (in bytes) of data stored in the file system.

Type: Long

Valid Range: Minimum value of 0.

Required: Yes

## MountTargetDescription

Provides a description of a mount target.

### Contents

#### FileSystemId

ID of the file system for which the mount target is intended.

Type: String

Required: Yes

#### IpAddress

Address at which the file system may be mounted via the mount target. Type: String

Required: No

#### LifeCycleState

Lifecycle state of the mount target.

Type: String

Valid Values: creating | available | deleting | deleted

Required: Yes

#### MountTargetId

System-assigned mount target ID.

Type: String

Required: Yes

#### NetworkInterfaceId

ID of the network interface that Amazon EFS created when it created the mount target.

Type: String

Required: No

#### Ownerld

AWS account ID that owns the resource.

Type: String

Required: No

#### SubnetId

ID of the mount target's subnet.

Type: String

Required: Yes

## Tag

A tag is a key-value pair. Allowed characters: letters, whitespace, and numbers, representable in UTF-8, and the following characters:  $+ - = . _ : /$ 

### Contents

#### Key

Tag key (String). The key can't start with aws:. Type: String Length Constraints: Minimum length of 1. Maximum length of 128. Required: Yes Value

Value of the tag key. Type: String Length Constraints: Maximum length of 256. Required: Yes

# **Document History**

The following table describes important changes to the Amazon Elastic File System User Guide.

- API version: 2015-02-01
- Latest documentation update: August 29, 2016

Change	Description	Date Changed
Increased tag support for file systems	Amazon EFS now supports 50 tags per file system. For more information on tags in Amazon EFS, see Managing File System Tags (p. 37).	In this release
General availability	Amazon EFS is now generally available to all users in the US East (N. Virginia), US West (Oregon), and EU (Ireland) regions.	June 28, 2016
File system limit increase	The number of Amazon EFS file systems that can be created per account per region increased from 5 to 10.	August 21, 2015
Updated Getting Started exercise	The Getting Started exercise has been updated to simplify the getting started process.	August 17, 2015
New guide	This is the first release of the Amazon Elastic File System User Guide.	May 26, 2015