
EFS-to-EFS Backup Solution Implementation Guide



EFS-to-EFS Backup Solution: Implementation Guide

Copyright © 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

Home	1
Overview	2
Cost	2
Architecture	2
Considerations	4
Incremental Backups	4
Consistent Backups	4
Sizing and Capacity	4
Burst Credits	5
Granular Backups	5
Encryption	5
Visualization	5
Cross-Account Backups	5
Solution Updates	5
Regional Deployment	5
Template	7
Deployment	8
Prerequisites	8
Launch the Stack	8
Restore a Backup	11
Security	14
Security Group	14
Resources	15
Appendix A: Instance Size	16
Appendix B: Logs	17
Appendix C: VPC Configuration	18
Appendix D: Large File Systems	19
Appendix E: Anonymous Data	20
Source Code	22
Revisions	23
.....	23

Automatically create incremental backups of an Amazon EFS file system on a customer-defined schedule

AWS Implementation Guide

AWS Solutions Builder Team

October 2017 (last update (p. 23): December 2019)

This implementation guide discusses architectural considerations and configuration steps for deploying the EFS-to-EFS backup solution on the Amazon Web Services (AWS) Cloud. It includes links to an [AWS CloudFormation](#) template that launches and configures the AWS services required to deploy this solution using AWS best practices for security and availability.

The guide is intended for IT infrastructure architects, administrators, and DevOps professionals who have practical experience architecting on the AWS Cloud.

Overview

Important

[AWS Backup](#), a fully managed backup service, enables you to centrally manage backups for Amazon EFS file systems. AWS recommends that you evaluate AWS Backup for your specific use case before you use this solution. AWS recommends using AWS Backup to implement a backup solution for Amazon EFS to recover from unintended user changes or deletions. For the most current Amazon EFS availability by Region, refer to the [Regional Product Services table](#). If AWS Backup is not available in your Region, use this EFS-to-EFS Backup solution.

Many Amazon Web Services (AWS) customers who use [Amazon Elastic File System](#) (Amazon EFS), a highly available and durable file storage service, choose to implement a backup solution to be able to recover from unintended changes or deletions to their file systems.

This guide provides infrastructure and configuration information for planning and deploying an EFS-to-EFS backup solution that automatically copies data from your Amazon EFS file system (the *source file system*) to another Amazon EFS file system (the *backup file system*).

The EFS-to-EFS backup solution automatically deploys the necessary AWS services, including Amazon CloudWatch and AWS Lambda, to create automated, incremental backups of Amazon EFS file systems on a customer-defined schedule. The solution is simple to configure and makes it easier to create and restore backups for data recovery and protection. You can use this solution to create backups for your file systems on a daily, weekly, or monthly basis and retain these backups as necessary to meet your business requirements.

Cost

You are responsible for the cost of the AWS services used while running this solution. The total cost of running this solution depends on the interval of the AWS Lambda function and the amount of storage your backup consumes. We recommend that you carefully consider your backup frequency and retention settings to avoid incurring unnecessary charges.

As of the date of publication, the total cost for running this solution with default settings in the US East (N. Virginia) Region is approximately **\$15.45 per month** plus variable data storage and Amazon Simple Notification Service (Amazon SNS) charges. The cost estimate assumes the following scenario:

- The solution creates a daily backup and an Amazon CloudWatch metrics dashboard
- The solution makes four AWS Lambda executions per backup process
- The solution uses five Amazon DynamoDB read capacity units and five write capacity units per month
- The Amazon Elastic Compute Cloud (Amazon EC2) instance runs for three hours daily

The cost to restore a backup depends on how long the restore Amazon EC2 instance runs (see [Restore a Backup \(p. 11\)](#)) and the number of Amazon SNS notifications the solution publishes. Prices are subject to change. For full details, see the pricing webpage for each AWS service you will be using in this solution.

Architecture Overview

Deploying this solution with the **default parameters** builds the following environment in the AWS Cloud.

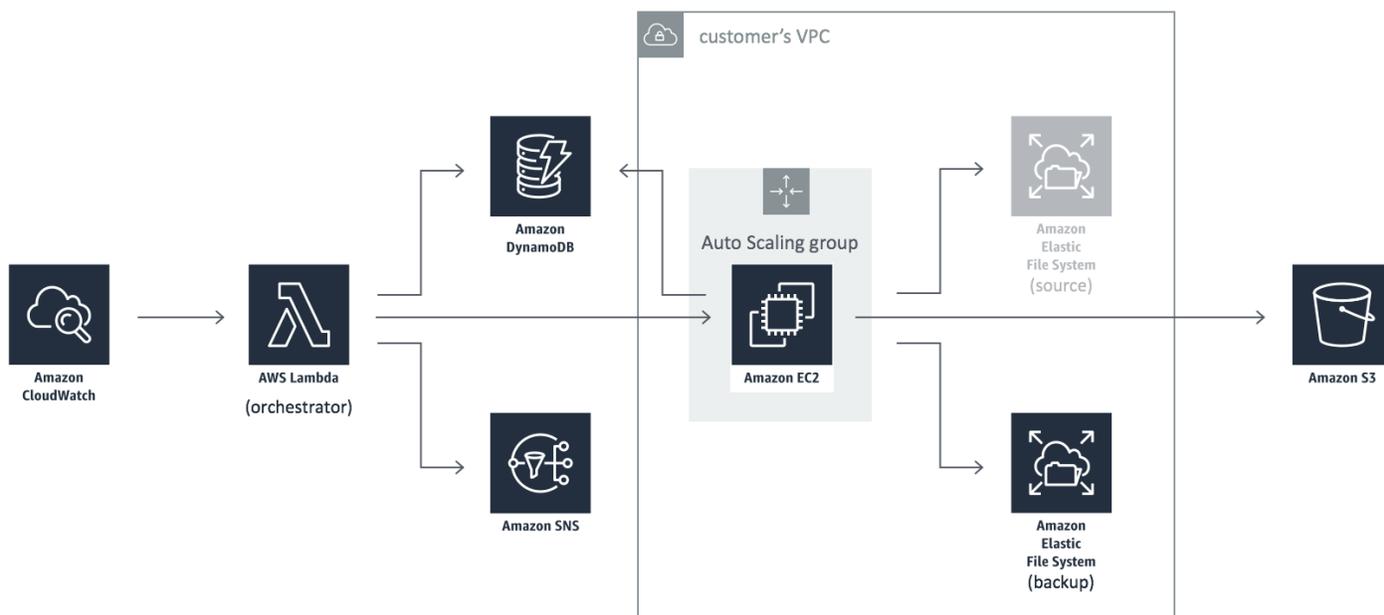


Figure 1: EFS-to-EFS backup solution architecture

The AWS CloudFormation template deploys two Amazon CloudWatch events, an AWS Lambda function, an Amazon DynamoDB table, an Amazon Simple Notification Service (Amazon SNS) topic, and an Amazon Simple Storage Service (Amazon S3) bucket. It also deploys an Amazon EC2 Auto Scaling group that launches and terminates an Amazon Elastic Compute Cloud (Amazon EC2) instance when a CloudWatch event occurs, and a backup Amazon Elastic File System (Amazon EFS) file system in your existing Amazon Virtual Private Cloud (Amazon VPC).

One Amazon CloudWatch event runs at an interval you specify during initial configuration. This event invokes the solution's *orchestrator* AWS Lambda function, which modifies the desired capacity of the Auto Scaling group to launch the Amazon EC2 instance, creates an ID for the backup, and stores details about the backup Amazon EFS file system in a DynamoDB table. The orchestrator function also creates a second CloudWatch event that stops the backup process if the backup process does not finish before a customer-defined amount of time (the backup window) has passed.

During boot, the Amazon EC2 instance mounts the source and backup Amazon EFS file systems and initiates the backup. When the backup process completes, the instance updates the backup activity details in the DynamoDB table, sends logs to the Amazon S3 bucket, and modifies the desired capacity of the Auto Scaling group to terminate the instance. If the instance cannot update the DynamoDB table, Amazon SNS sends a notification to subscribed email addresses.

If the backup window expires before the backup process completes, the second CloudWatch event invokes the orchestrator function to update the desired capacity of the Auto Scaling group to terminate the instance.

If the backup is unsuccessful, Amazon SNS sends a notification to subscribed email addresses. You can also choose to be notified if the backup is successful.

Design Considerations

Incremental Backups

This solution captures the state of an Amazon EFS file system at a point in time. If you specify a backup window long enough to copy your entire file system, the solution will copy the entire file system or the entire subdirectory the first time it creates a backup. When the solution creates future backups for that file system, it copies only the files and directories that have changed, or been added or removed since the last backup.

We recommend that you launch the solution for the first time with [a large instance type \(p. 16\)](#), and a backup window large enough to copy your entire file system. After the first backup completes, update the running stack with a smaller instance type and backup window to reduce costs and save burst credits.

Consistent Backups

This solution uses *fpsync*, a tool that synchronizes directories in parallel using *fpart* (sorts and packs files into partitions) and *rsync* (a file-copying tool), to copy the source file system to the backup file system. Note that this solution might exclude any data written while *fpsync* or the backup process are running. To ensure consistent backups, we recommend that you do not perform writes on the source Amazon EFS file system for the duration of the backup process. Note that after the solution creates the initial backup, future backups are incremental and will take less time to complete.

Sizing and Capacity

This solution uses a single Amazon Elastic Compute Cloud (Amazon EC2) instance (c5.xlarge). The maximum throughput you can drive per NFS client on an Amazon EC2 instance is 250 MB/s. All Amazon EFS 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. The size of the file system determines the portion of time a file system can burst. Amazon EFS uses a credit system to determine when file systems can burst. For more information about how the credit system works, see [Amazon EFS Performance](#) in the *Amazon EFS User Guide*.

Large file systems might cause this solution to hit the Amazon EC2 instance throughput limit. Customers who want to back up large file systems can launch multiple deployments of the solution with different source prefixes that point to different locations in the source file system. For example, a customer with a large Amazon EFS file system that contains a `home` directory and an `appdata` directory can deploy two solution stacks: one that points to the `home` directory (`<efs-mount-point>:/home`) and one that points to the `appdata` directory (`<efs-mount-point>:/appdata`). Customers can also launch multiple deployments of the solution to back up multiple Amazon EFS file systems in an AWS Region.

We recommend that you use a rigorous performance testing and optimization process to choose the right instance type for your use case. For more information, see [Appendix A \(p. 16\)](#).

File systems with a large number of files (around one million or more) can also cause the backup process to fail. For more information about file systems with a large number of files, see [Appendix D \(p. 19\)](#).

Burst Credits

This solution will consume burst credits while creating backups, which could impact your production workload. We recommend that you verify that you have sufficient burst credits available before you start the backup process. You can also change the solution's default Amazon EC2 instance type to change how the solution consumes burst credits. For more information, see [Appendix A \(p. 16\)](#).

Granular Backups

The EFS-to-EFS backup solution enables customers to specify any valid directory from their source file system. To back up part of an Amazon EFS file system, customers can specify the applicable subdirectory as the source, and the solution will copy only files and directories from that subdirectory. Continuing from the previous example (see [Sizing and Capacity \(p. 4\)](#)), if a customer specifies the `appdata` directory as the source, the solution will copy only the `appdata` directory to the backup file system.

Encryption

By default, this solution leverages Amazon EFS encryption so you can encrypt your backups at rest. Amazon EFS integrates with AWS Key Management Service (AWS KMS) customer master keys (CMKs), and uses an industry-standard AES-256 encryption algorithm to encrypt EFS data and metadata. For more information, see [Cryptography Basics](#) in the *AWS Key Management Service Developer Guide*.

Visualization

The EFS-to-EFS backup solution includes optional dashboards that allow you to visualize Amazon EFS I/O data for both the source and backup file systems during the backup and restore processes. To view the dashboards, open the Amazon CloudWatch console and select **Dashboards**.

Cross-Account Backups

This solution will not back up Amazon EFS file systems that are mounted from a different account or VPC by default. To back up these types of file systems, you must modify the solution. For more information, see [Mounting EFS File systems from Another Account or VPC](#) in the *Amazon EFS User Guide*.

Solution Updates

EFS-to-EFS backup version 1.5 uses the most up-to-date Node.js runtime. Version 1.4 uses the Node.js 8.10 runtime, which reaches end-of-life on December 31, 2019. In January, AWS Lambda will block the create operation and, in February, Lambda will block the update operation. For more information, see [Runtime Support Policy](#) in the *AWS Lambda Developer Guide*.

To continue using this solution with the latest features and improvements, you can update the stack.

Regional Deployment

This solution uses the Amazon EFS service, which is currently available in specific AWS Regions only. Therefore, you must launch this solution in an AWS Region where Amazon EFS is available. (For the

most current Amazon EFS availability by region, see <https://aws.amazon.com/about-aws/global-infrastructure/regional-product-services/>.) Also, you must deploy this solution in the same AWS Region as your source Amazon EFS file system. You can launch multiple deployments of the solution in a single AWS Region to back up multiple Amazon EFS file systems in that region.

AWS CloudFormation Template

This solution uses AWS CloudFormation to automate the deployment of the EFS-to-EFS backup solution. It includes the following AWS CloudFormation template, which you can download before deployment:

[View
Template](#)

efs-to-efs-backup.template: Use this template to launch the EFS-to-EFS backup solution and all associated components. The default configuration deploys two Amazon CloudWatch events, an AWS Lambda function, an Amazon DynamoDB table, an Amazon Simple Notification Service (Amazon SNS) topic, an Amazon Simple Storage Service (Amazon S3) bucket, an Amazon EC2 Auto Scaling group that launches and terminates an Amazon Elastic Compute Cloud (Amazon EC2) instance, and an Amazon Elastic File System (Amazon EFS) file system to store backups.

Automated Deployment

Before you launch the automated deployment, please review the considerations and prerequisites discussed in this guide. Follow the step-by-step instructions in this section to configure and deploy the EFS-to-EFS backup solution into your account.

Time to deploy: Approximately five minutes

Prerequisites

Before you start, you must have a source Amazon Elastic File System (Amazon EFS) file system with mount targets in an Amazon Virtual Private Cloud (Amazon VPC) network. The VPC network must have the following:

- At least two subnets in different Availability Zones (AZs)
- Public subnets or private subnets with a route to a network address translation (NAT) gateway. For more information, see [Appendix C \(p. 18\)](#).
- A set of DHCP options configured to use the DNS server provided by Amazon. For more information about the Amazon DNS server, see [DHCP Options Sets](#) in the *Amazon VPC User Guide*.
- DNS hostnames enabled. For more information, see [Viewing DNS Hostnames for Your EC2 Instance](#) in the *Amazon VPC User Guide*.
- The source file system mount targets, the Amazon EC2 Auto Scaling group, and the destination file system mount targets in the same AZs
- Security group rules that grant the solution's Amazon EC2 instance inbound access to your source Amazon EFS file system through the mount target using the Network File System (NFS) port. In the source EFS mount target security group, allow inbound access for the NFS port to the VPC Classless Inter-Domain Routing (CIDR) block or specific subnet CIDR blocks.

Launch the Stack

This automated AWS CloudFormation template deploys the EFS-to-EFS backup solution.

Note

You are responsible for the cost of the AWS services used while running this solution. See the [Cost \(p. 2\)](#) section for more details. For full details, see the pricing webpage for each AWS service you will be using in this solution.

1. Sign in to the AWS Management Console and click the button below to launch the `efs-to-efs-backup` AWS CloudFormation template.



You can also [download the template](#) as a starting point for your own implementation.

- The template is launched in the US East (N. Virginia) Region by default. To launch the solution in a different AWS Region, use the region selector in the console navigation bar.

Note

This solution uses the Amazon EFS service, which is currently available in specific AWS Regions only. Therefore, you must launch this solution in an AWS Region where Amazon EFS is available. For the most current Amazon EFS availability by region, see <https://aws.amazon.com/about-aws/global-infrastructure/regional-product-services/>. Also, you must deploy this solution in the same AWS Region as your source Amazon EFS file system.

- On the **Select Template** page, verify that you selected the correct template and choose **Next**.
- On the **Specify Details** page, specify a name for your solution stack.
- Under **Parameters**, review the parameters for the template, and modify them as necessary. This solution uses the following default values.

Parameter	Default	Description
Source EFS	<Requires input>	The source Amazon EFS file system ID
Interval Label	daily	The tag that identifies your backups. Choose the applicable frequency: daily, weekly, or monthly.
Retain	7	The number of backups you want to retain
Folder Label	efs-backup	The folder name for your backups
Backup Window	180	The maximum amount of time in minutes that the backup process has to complete
Backup Schedule	cron(0 2 * * ? *)	Enter the scheduled expression (Cron syntax) that specifies when to run the CloudWatch event. All times are in UTC in 24-hour format. For example, cron(0 2 * * ? *) to run the backup at 2 a.m. UTC daily.
Backup Prefix	/	The source prefix for the backup. To back up part of the source file system, specify the applicable subdirectory. For example, enter /appdata to back up the appdata directory.
EFS Mode	generalPurpose	The performance mode for the backup Amazon EFS file system
Subnet IDs	<Requires input>	Comma-delimited list of two subnet IDs for the Auto Scaling

Parameter	Default	Description
		group. The backup mount targets will be created in these subnets. Note These subnets must be in the same Amazon VPC and Availability Zones as the source file system mount targets, and they must have outbound internet access.
VPC ID	<Requires input>	The virtual private cloud (VPC) ID where the source and backup mount targets are located
Email	<Requires input>	The email address to subscribe to Amazon SNS notifications
Success Notification	Yes	Choose whether to be notified when backups are created successfully. Note The solution always sends notifications when backups fail.
Dashboard	Yes	Choose whether to create a dashboard for metrics
EFS Encryption	Yes	Choose whether to encrypt the backup Amazon EFS file system

6. Choose **Next**.
7. On the **Options** page, choose **Next**.
8. On the **Review** page, review and confirm the settings. Select the checkbox acknowledging that the template will create AWS Identity and Access Management (IAM) resources.
9. Choose **Create** to deploy the stack.

You can view the status of the stack in the AWS CloudFormation console in the **Status** column. You should see a status of **CREATE_COMPLETE** in approximately five minutes.

Note

In addition to the primary AWS Lambda function, this solution includes the `solution-helper` and `AMIInfoFunction` Lambda functions, which run only during initial configuration or when resources are updated or deleted.

When running this solution, you will see these Lambda functions in the AWS console, but only the primary function is regularly active. However, do not delete the `solution-helper` and `AMIInfoFunction` functions as they are necessary to manage associated resources.

Restore a Backup

When the EFS-to-EFS backup solution creates a backup, the solution labels the backup with the **Interval Label** you specify during deployment and a number to identify the backup. For example, if you choose `daily` as your **Interval Label**, your backups will be named `daily.0`, `daily.1`, `daily.2`, etc. The most recent backup will be `daily.0`.

To restore an Amazon Elastic File System (Amazon EFS) file system, identify the **Interval Label** and the number of the backup you want to restore and complete the following procedure.

1. Sign in to the AWS Management Console and click the button below to launch the `efs-to-efs-restore` AWS CloudFormation template.



You can also [download the template](#) as a starting point for your own implementation.

2. The template is launched in the US East (N. Virginia) Region by default. To launch the solution in a different AWS Region, use the region selector in the console navigation bar.
3. On the **Select Template** page, verify that you selected the correct template and choose **Next**.
4. On the **Specify Details** page, specify a name for your solution stack.
5. Under **Parameters**, review the parameters for the template, and modify them as necessary. This solution uses the following default values.

Parameter	Default	Description
Source EFS	<i><Requires input></i>	The source Amazon EFS file system ID
Backup EFS	<i><Requires input></i>	The backup Amazon EFS file system ID
Interval Label	<code>daily</code>	The tag that identifies the backup you want to restore. The tag must match the tag you chose for Interval Label during deployment. For example, if you chose <code>daily</code> during deployment, choose <code>daily</code> for this parameter.
Backup Number	<code>0</code>	The number of the backup you want to restore. The most recent backup is <code>0</code> .
Folder Label	<code>efs-backup</code>	The folder name for your backups Note The folder name must match the folder

Parameter	Default	Description
		name you specified in the Folder Label parameter for the <code>efs-to-efs-backup</code> template.
Restore Prefix	/	The source prefix for the restored file system Note If you backed up a subdirectory, enter the subdirectory for this parameter. For example, enter /appdata for your appdata backup.
Restore Subdirectory	/	The subdirectory you want to restore. For example, /<directory-name>/ Note You must enter the trailing "/".
Restore Log Bucket	<Requires input>	The Amazon S3 bucket to store logs. We recommend you use the bucket that stores your backup logs.
Subnet IDs	<Requires input>	Comma-delimited list of two subnet IDs for the Auto Scaling group. Note The subnets must be in the same Availability Zones as the source and backup file system mount targets.
VPC ID	<Requires input>	The virtual private cloud (VPC) ID where the source and backup Amazon EFS file systems are located
Security Group ID	<Requires input>	The security group ID of an existing Amazon EC2 security group in your VPC to provide access to your existing Amazon EFS file system
Email	<Requires input>	The email address to subscribe to Amazon SNS notifications

Parameter	Default	Description
Dashboard	Yes	Choose whether to create a dashboard for metrics

6. Choose **Next**.
7. On the **Options** page, choose **Next**.
8. On the **Review** page, review and confirm the settings. Be sure to check the box acknowledging that the template will create AWS Identity and Access Management (IAM) resources.
9. Choose **Create** to deploy the stack.

You can view the status of the stack in the AWS CloudFormation console in the **Status** column. You should see a status of **CREATE_COMPLETE** in approximately five minutes.

When your Amazon EFS file system is restored to the backup, the solution sends an Amazon Simple Notification Service (Amazon SNS) notification.

Security

When you build systems on AWS infrastructure, security responsibilities are shared between you and AWS. This shared model can reduce your operational burden as AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate. For more information about security on AWS, visit the [AWS Security Center](#).

Security Group

The security group created in this solution is designed to allow all traffic to flow from the Amazon Elastic Compute Cloud (Amazon EC2) instance to the backup Amazon Elastic File System (Amazon EFS) file system within the same security group. The Amazon EC2 instance must have inbound access to the source Amazon EFS file system. We recommend that you review the security groups and further restrict access as needed once the deployment is up and running. For more information, see [Security Groups for Amazon EC2 Instances and Mount Targets](#) in the *Amazon EFS User Guide*.

Additional Resources

AWS services documentation

- [Amazon Elastic File System](#)
- [AWS CloudFormation](#)
- [Amazon CloudWatch](#)
- [AWS Lambda](#)
- [Amazon DynamoDB](#)
- [Amazon Simple Notification Service](#)

Appendix A: Amazon EC2 Instance Size

Amazon Elastic Compute Cloud (Amazon EC2) provides a wide selection of instance types and sizes optimized to fit different use cases, giving you the flexibility to choose the appropriate instance type for your workload.

By default, the EFS-to-EFS backup solution uses the `c5.xlarge` Amazon EC2 instance type, but you can choose a different instance type depending on how you want the solution's file system to consume burst credits. For example, a customer with on-going production reads to their source Amazon Elastic File System (Amazon EFS) file system might choose a smaller instance type for the solution to save burst credits for the production load.

To change the instance size, modify the following mapping in the AWS CloudFormation template:

```
c5.xlarge: {"Arch": "HVM64"}
us-east-1: {"InstanceSize": "c5.xlarge"}
```

For example, to change the instance size to a `c5.large`, make the following change:

```
c5.large: {"Arch": "HVM64"}
us-east-1: {"InstanceSize": "c5.large"}
```

We recommend that you use a compute-optimized instance type for optimized performance and cost efficiency. If you choose to use a different instance type, we recommend that you use a rigorous performance testing and optimization process to choose an instance type that optimizes the consumption of burst credits. For more information about instance sizes, see [Amazon EC2 Instance Types](#).

Appendix B: Logs

The EFS-to-EFS backup solution sends logs to an Amazon Simple Storage Service (Amazon S3) bucket. The following prefixes are used in the Amazon S3 bucket:

- **ec2-logs:** Contains the `cloud-init` log and detailed output of the `efs-backup-fpsync.sh` script.
- **efs-backup-log:** Contains the `fpsync` and `rsync` process logs for the `efs-to-efs-backup` template.
- **ssm-logs:** Contains the `stdout` and `stderr` logs for the `ssm.sh` script.
- **efs-restore-logs:** Contains the logs from the restore process after the files are uploaded.

Appendix C: VPC/Subnet Configuration

Before you start, review the architecture of your source Amazon Elastic File System (Amazon EFS) file system. It should look similar to the following example diagram.

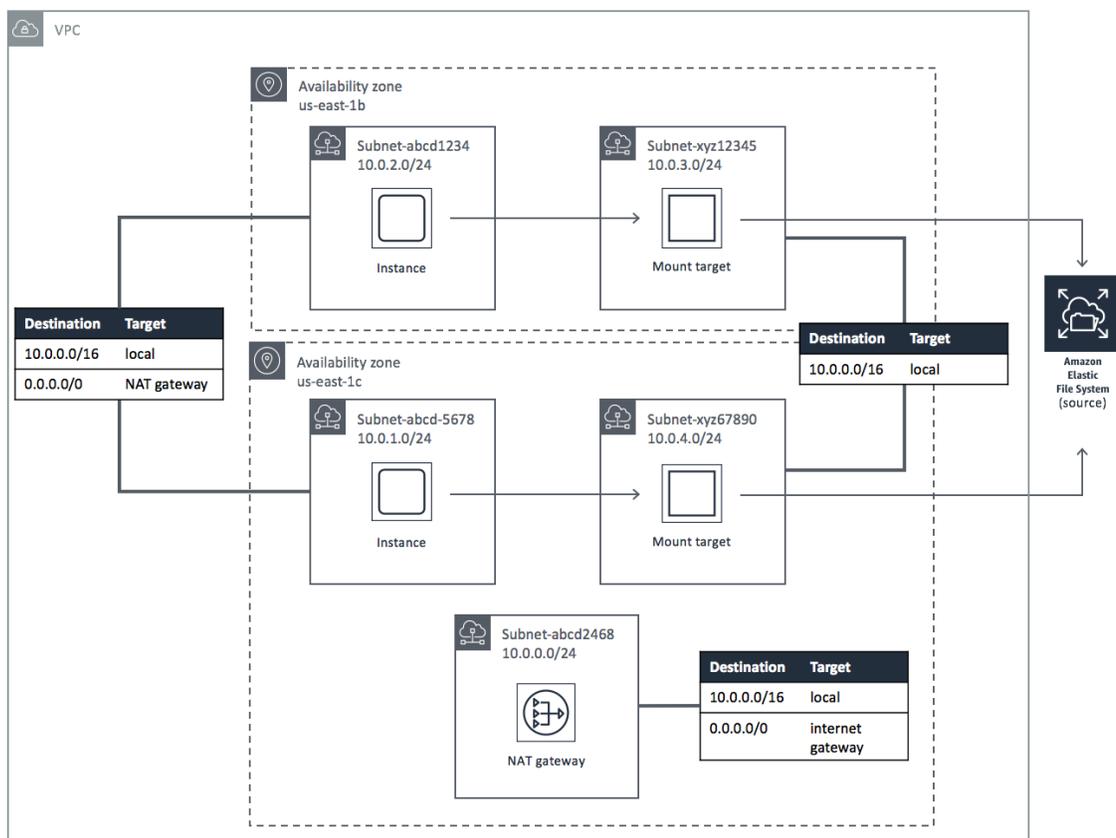


Figure 2: EFS-to-EFS backup solution VPC/subnet configuration

Appendix D: Large File Systems

The EFS-to-EFS backup solution uses *fpsync*, a tool that synchronizes directories in parallel using *fpart* (sorts and packs files into partitions) and *rsync* (a file-copying tool), to copy the source file system to the backup file system. File systems with a large number of files (around one million files or more) can cause the *fpart* process to fail. If the process fails, the solution will generate an incomplete backup error message.

```
"InstanceType": "c5.xlarge",
"IntervalTag": "daily",
"Message": "The EFS backup was incomplete. Either backup window expired before full backup or fpsync process was not completed."
```

Figure 3: Sample incomplete backup process error message

You can also review the *fpsync* output logs in your Amazon Simple Storage Service (Amazon S3) bucket to determine if the *fpsync* synchronization ran successfully. The following log message is an example of a successful *fpsync* synchronization:

```
which: no mail in (/sbin:/usr/sbin:/bin:/usr/bin:/usr/local/bin)
====> Job name: 1548270480-24315
====> Analyzing filesystem...
====> Waiting for sync jobs to complete...
<==== Parts done: 1/1 (100%), remaining: 0
<==== Fpsync completed without error.
```

Figure 4: Sample successful fpsync synchronization message

If the *fpsync* synchronization is not successful, the solution will generate an incomplete log message.

```
which: no mail in (/sbin:/usr/sbin:/bin:/usr/bin:/usr/local/bin)
====> Job name: 1548273158-24438
====> Analyzing filesystem...
```

Figure 5: Sample incomplete fpsync synchronization message

For customers with Amazon EFS file systems with a large number of files, we recommend that you launch multiple deployments of the solution with different source prefixes that point to different locations in your source file system. For more information, see [Sizing and Capacity \(p. 4\)](#).

Appendix E: Collection of Operational Metrics

This solution includes an option to send anonymous operational metrics to AWS. We use this data to better understand how customers use this solution and related services and products. When enabled, the following information is collected and sent to AWS:

- **Solution ID:** The AWS solution identifier
- **Unique ID (UUID):** Randomly generated, unique identifier for each solution deployment
- **Timestamp:** Data-collection timestamp
- **Backup ID:** The interval label and number
- **Backup Start Time:** The time the backup starts
- **Backup Stop Time:** The time the backup stops
- **Backup Window:** The amount of time the backup process has to complete
- **Source EFS Size:** The size of the source Amazon Elastic File System (Amazon EFS) file system
- **Destination EFS Size:** The size of the backup Amazon EFS file system
- **Instance Type:** The Amazon Elastic Compute Cloud (Amazon EC2) instance type
- **Retain:** The number of backups you retain
- **S3 Bucket Size:** The Amazon Simple Storage Service (Amazon S3) bucket size
- **Source Burst Credit Balance:** The burst credit balance of your source Amazon EFS file system
- **Source Burst Credit Balance Post Backup:** The burst credit balance of your source Amazon EFS file system after the backup process completes
- **Source Performance Mode:** The performance mode of your source Amazon EFS file system
- **Destination Performance Mode:** The performance mode of your backup Amazon EFS file system
- **Number of Files:** The number of files in your source Amazon EFS file system that are analyzed for back up by the fpsync process
- **Number of Files Transferred:** The number of files transferred from the source Amazon EFS file system to the backup file system
- **Total File Size:** The size (in bytes) of your Amazon EFS file system that are analyzed for back up by the fpsync process
- **Total Transferred File Size:** The size (in bytes) of the files transferred from the source Amazon EFS file system to the backup file system
- **Region:** The AWS Region with your source Amazon EFS file system
- **Create Hard Links Start Time:** The time the process to create hard links for the most recent backup starts
- **Create Hard Links Stop Time:** The time the process to create hard links for the most recent backup stops
- **Remove Snapshot Start Time:** The time the process to delete the oldest backup starts
- **Remove Snapshot Stop Time:** The time the process to delete the oldest backup stops
- **Rsync Delete Start Time:** The time the rsync delete process starts
- **Rsync Delete Stop Time:** The time the rsync delete process stops

Note that AWS will own the data gathered via this survey. Data collection will be subject to the [AWS Privacy Policy](#). To opt out of this feature, modify the AWS CloudFormation template mapping section as follows:

```
Mappings:  
  Map:  
    send-data: {"SendAnonymousData": "Yes"},
```

to

```
Mappings:  
  Map:  
    send-data: {"SendAnonymousData": "No"},
```

Source Code

You can visit our [GitHub repository](#) to download the templates and scripts for this solution, and to share your customizations with others.

Document Revisions

Date	Change
October 2017	Initial release
March 2018	Added information on new Security Group ID parameter
May 2018	Added information on new Restore Subdirectory parameter
July 2018	Added information on how to change the instance type with template mapping
February 2019	Added recommendations for file systems with a large number of files
December 2019	Added information on support for Node.js update

Notices

This implementation guide is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

The EFS-to-EFS backup solution is licensed under Apache License Version 2.0 available at <https://www.apache.org/licenses/LICENSE-2.0>