

Load, Store, and Protect Linux-Based NFS Workloads in AWS

A Cloud Migration Guide with AWS DataSync, Amazon Elastic File System, and AWS Backup

September 2019



Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. 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.

© 2019 Amazon Web Services, Inc. or its affiliates. All rights reserved.

Contents

Introduction	1
Service Overview.....	2
AWS DataSync.....	2
Amazon Elastic File System	2
AWS Backup	3
Migration Guide	3
Prerequisites.....	3
Setup and Configuration	4
Operations and Monitoring.....	8
Pricing and Cost Considerations	9
Conclusion	11
Contributors	11
Document Revisions.....	11

Abstract

This paper outlines a strategy for migrating Linux-based applications that require Network File System (NFS) storage to AWS. For Windows-based applications, we offer [Amazon FSx for Windows File Server](#). Using this paper, learn how to deploy AWS DataSync to quickly and effectively transfer files from an on-premises file system to Amazon Elastic File System, how to use the Amazon Elastic File System Infrequent Access storage class and file lifecycle management to reduce storage costs by up to 92%, and how to use AWS Backup for automated, centrally managed data protection and recovery.

Introduction

The depth and breadth of functionality available through cloud computing continues to grow. Organizations are adopting cloud services to gain additional business agility and lower their IT total cost of ownership. However, it can be challenging to keep up with the latest available options and to develop a plan for building new applications or moving existing ones into the cloud with the right architecture.

This whitepaper focuses on best practices for migrating Linux file-based applications to the cloud. The vast majority of data that exists today is in file format on file system storage. When considering cloud file storage options, you need to satisfy business objectives such as high availability, security and compliance, and backup and recovery, while managing cost and planning for seamless migration.

AWS provides simple and cost-effective options for loading, storing, and protecting your file-based data in the cloud. To load data, you can use AWS DataSync, a data transfer service that makes it easy for you to automate moving data between on-premises storage and the cloud. DataSync automatically handles many of the tasks related to data transfers, eliminating the need to write scripts or modify your applications. You can use DataSync to copy data to AWS using AWS Direct Connect or the internet, at speeds up to 10 times faster than open-source tools.

To store Linux-based file data natively in a file system, you can use Amazon Elastic File System (Amazon EFS). For Windows-based file data, we offer [Amazon FSx for Windows File Server](#). This paper focuses on Linux-based workloads.

Amazon EFS provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. In addition to the data protection features provided by Amazon EFS, such as high availability, high durability, and encryption, you also can use AWS Backup, a fully-managed backup service that makes it easy to centralize and automate file system backups. With just a few clicks in the AWS Backup console, you can create backup policies that automate backup schedules and retention management.

This whitepaper discusses these services in detail, highlighting features, configurability, and cost, and shows how to use them together as part of an overall migration plan.

Service Overview

AWS DataSync

AWS DataSync is a managed data transfer service that simplifies moving, copying, and synchronizing large amounts of data between on-premises storage systems and AWS storage services. This whitepaper focuses on using DataSync in conjunction with Amazon Elastic File System (Amazon EFS) in particular.

DataSync accelerates and automates data transfer—removing the need to modify applications, develop scripts, or manage infrastructure. DataSync uses a purpose-built protocol to accelerate and secure data transfer at speeds up to 10 times faster than open-source tools.

DataSync has simple, per gigabyte pricing, so that you pay only for what you use. DataSync uses a software agent to handle transferring files and supports both one-time migrations and periodic synchronizations of data between on-premises file systems and EFS.

In terms of data protection, DataSync provides automatic encryption of data in transit and validation of data integrity.

Amazon Elastic File System

Amazon Elastic File System (Amazon EFS) provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. It is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files, eliminating the need to provision and manage capacity to accommodate growth.

Amazon EFS offers two storage classes: the Standard storage class and the [Infrequent Access storage class](#) (EFS IA). EFS IA provides price/performance that's cost-optimized for files not accessed daily. By enabling EFS Lifecycle Management on your file system, files not accessed according to the lifecycle policy you choose will be automatically and transparently moved into EFS IA. EFS IA provides storage cost savings of up to 92%.

Amazon EFS is well suited to support a broad spectrum of use cases from home directories to business-critical applications. Customers can use EFS to lift-and-shift existing enterprise applications to the AWS Cloud. Other use cases include: big data

analytics, web serving and content management, application development and testing, media and entertainment workflows, database backups, and container storage.

EFS supports sharing storage across up to thousands of Amazon Elastic Compute Cloud (Amazon EC2) instances. You can access Amazon EFS using a Virtual Private Cloud (VPC), over Virtual Private Networks (VPN), or from an on-premises datacenter using AWS Direct Connect.

AWS Backup

AWS Backup is a fully managed backup service that makes it easy to centralize and automate the back up of data across AWS services in the cloud. AWS Backup provides a policy-based backup solution that simplifies your backup management and enables you to meet your business and regulatory backup compliance requirements. Using AWS Backup, you can backup and restore your Amazon EFS file systems natively.

Migration Guide

To get the most out of this whitepaper, we recommend that you have basic familiarity with the setup and use of each service: AWS DataSync, Amazon Elastic File System, and AWS Backup. The following guide provides considerations for migrating an on-premises file system to AWS, including configuration options.

Prerequisites

To execute the steps outlined in this whitepaper, it is assumed that you have an active AWS account with the permissions required to create new resources for each of the mentioned services. Note that you are responsible for any charges incurred from creating and using these resources.

To transfer files from on-premises to AWS using DataSync, you need a virtualization system, such as VMware ESXi, that is capable of running the DataSync agent. You will also need access to a web browser to activate the DataSync agent. For more information, see the [DataSync documentation](#).

Setup and Configuration

Amazon Elastic File System

Create an Amazon EFS file system using the AWS Management Console, Command Line Interface (CLI), or Software Development Kit (SDK). You can use this file system to store the data that you migrate from your on-premises environment. As a fully managed, elastic file system, there is no setup or provisioning needed and no need to plan for or manage storage capacity. EFS file systems automatically grow and shrink as you add and remove files.

General

Amazon EFS is designed to be highly available and highly durable. EFS offers a [Service Level Agreement \(SLA\)](#) for 3 9's of availability, and is designed for 11 9's of data durability. To achieve this level of availability and durability, EFS replicates data within and across multiple Availability Zones (AZs). EFS is designed to tolerate the unavailability of a single AZ without sacrificing performance or features.

Amazon EFS file systems can be accessed by thousands of clients simultaneously across multiple AZs. To access your file system, you create highly available resources called mount targets, one per AZ, and specify which VPC Security Groups are allowed to access them. The easiest way for your clients to access your file system is by using its DNS name. EFS file system DNS names automatically resolve to the IP address of the mount target located in the same AZ as the client making the request—optimizing performance and removing any potential cross-AZ network charges.

When creating a file system, you can select some configuration options. For data protection and to satisfy compliance requirements, we recommend enabling encryption of data at rest. Enabling encryption of data at rest must be done when you create your file system.

Performance

You also can choose a performance mode and a throughput mode for your file system. EFS offers two performance modes: General Purpose and Max I/O. Max I/O scales to higher amounts of aggregate throughput than General Purpose mode, but also has slightly higher metadata operation latencies. Like enabling encryption, you must select your file system performance mode when creating your file system. For the purposes of this whitepaper, we recommend selecting the default performance mode, General

Purpose. Additional guidance is provided later in this whitepaper for evaluating which performance mode is optimal for your file system.

EFS offers two throughput modes: Bursting and Provisioned. In the default Bursting throughput mode, EFS file system performance scales linearly with the size of the file system's Standard storage class. When you create an Amazon EFS file system in Bursting Throughput mode, you get an allocation of 2.1 TB of burst credits. All EFS file systems are able to burst up to 100 MB/s of throughput when using Bursting Throughput mode, with file systems larger than 1 TiB able to burst to twice their baseline throughput. Consider using EFS Provisioned Throughput if you know how much throughput your applications require, and if it's higher than the amount that your storage provides. For more information, refer to the whitepaper [Amazon Elastic File System: Choosing Between the Different Throughput & Performance Modes](#).

In the next section, we will discuss AWS DataSync and how you can use it to quickly transfer data into EFS. If you're planning to migrate more than 2.1 TB of data to a new EFS file system, we recommend provisioning at least 100 MB/s of throughput which will prevent performance slowdowns if you run out of EFS burst credits. If you must transfer data to an empty EFS file system at speeds greater than 100 MB/s, you will need to use Provisioned Throughput.

Storage Classes and Lifecycle Management

Amazon EFS offers two storage classes: Standard and Infrequent Access. The EFS Infrequent Access (EFS IA) storage class is intended for files that are not accessed every day and stores them on slower storage media at a lower cost. You can start using EFS IA by enabling Lifecycle Management for your file system. With Lifecycle Management enabled, EFS automatically moves files that have not been accessed according to the policy you choose from the Standard storage class to the EFS IA storage class, which can reduce your storage costs by up to 92%. Each EFS file system transparently serves files from both storage classes. We recommend enabling Lifecycle Management to automatically save cost as your file access patterns change and files are less frequently accessed.

AWS DataSync

To perform one-time data migrations or periodic synchronizations from an on-premises file system to Amazon EFS, use AWS DataSync. This whitepaper assumes that you have an AWS Direct Connect connection, but Direct Connect is not required to use DataSync. Using Direct Connect, you can transfer at speeds up to 10 Gbps using

DataSync. At 10 Gbps sustained, you can transfer 1 terabyte of data from on-premises to the cloud in around 13 minutes, or a petabyte in around 9.5 days.

For data protection, AWS DataSync automatically encrypts data in transit and performs data integrity checks. DataSync uses simple, per-GB pricing. DataSync transfers are always incremental and DataSync has built-in compression and de-duplication to minimize the amount of network activity needed to get the job done.

To get started, you first download and activate the DataSync agent as a virtual machine in your on-premises environment. The DataSync agent handles the actual data transfer between your on-premises file system and EFS. After you activate the agent, proceed to the DataSync management console to create a 'task', which is the complete definition of a data transfer that will handle the transfer between two locations: an on-premises Network File System (NFS) server and Amazon EFS. See the [AWS DataSync Getting Started Guide](#) for detailed step-by-step instructions. We recommend using NFS version 4.1 (or later) for the best performance, but versions 4.0 and 3 of the NFS protocol are also supported.

DataSync provides a number of configuration options that you can set to your needs. First, make sure that you have set the proper network configuration, including firewall and VPC security group rules. You will need to enable NFS traffic (port 2049) in your on-premises firewall and AWS VPC, and you will want to make sure that DataSync is configured to belong to at least one security group that is shared with your EFS mount targets.

DataSync offers options on how to treat your file system metadata. By default, DataSync will preserve ownership information, permissions, and timestamps of your on-premises file system once it's been moved to EFS. We recommend keeping these default settings in place when performing your migration.

DataSync also allows you to turn off the automatic data verification step (i.e., consistency check) it performs after the transfer is complete. Data verification can take a significant amount of time depending on how many files are in your file system and the rate of change to those files. If you are able to stop any file system activity against your on-premises file system before and for the duration of the transfer, you can disable data verification to save time.

DataSync provides a few options for increasing or decreasing file transfer bandwidth and performance. One option is to set a bandwidth limit for your transfer tasks—you can choose to optionally set this limit to enforce a quality of service policy to prevent DataSync from consuming too many network resources. Alternatively, if you want

DataSync to fully saturate the network and load data into EFS as quickly as possible, you can deploy and allocate as many as four DataSync agents to a given task. You may choose this option if you have millions of files in your source file system. You also can use multiple agents when you have network limitations on-premises (e.g., when not using Direct Connect) and want to get the most out of the limited network resources available.

After your DataSync task is set up, execute it on-demand or periodically using the DataSync API to incrementally copy changes from your on-premises NFS share to EFS.

AWS Backup

For additional data protection after you've copied your data to your EFS file system, you might want to back up your file system using AWS Backup. AWS Backup allows you to create scheduled backups to satisfy compliance requirements and recover from unintended events, such as the accidental deletion of files from your file system. Each backup is incremental—only changes to the file system are recorded.

With AWS Backup, you issue a backup by creating a resource called a backup plan. You can choose from some predefined backup plans that AWS Backup provides, or you can build a new one from scratch. An example of a predefined backup plan is one that executes daily and has a 35-day retention policy.

With your backup plan, you can select a frequency and schedule, or you can execute it one time on-demand. Optionally, you can choose when your backup data migrates to a lower-cost 'cold' storage, or you can automatically purge older backups using an expiration policy. When using cold storage for backups, there is an additional retrieval fee if you need to restore it, and it takes a few extra hours to retrieve the backup from cold storage.

After you create your backup plan, you can choose which of your EFS file systems are governed by it. You can do this in one of two ways. One way is by using a tag-based mechanism: you can have your backup plan automatically back up any file system tagged with a user-defined string such as "Production". Tag-based backup plans will cover file systems created in the future that are tagged without requiring any changes to the backup plan. The other way to back up a file system is to specify it by its ID in the backup plan.

With AWS Backup, you create resources called backup vaults to organize and store your backups. When you create a backup vault, you select a KMS key to encrypt your backups. When you back up an EFS file system, it is always stored encrypted.

You can easily restore your file system from backup using the AWS Backup console. You can choose to restore the backup to a specific folder in an existing file system, or you can choose to restore the backup to a brand new file system.

Operations and Monitoring

Now that you've set up your EFS file system, your DataSync agent and task, and your backup plan, you are ready to move files from on-premises to the cloud and begin to run applications against EFS. AWS provides tools and metrics to help you evaluate and optimize performance on an ongoing basis, both from within each respective AWS service as well as through dedicated operational services such as Amazon CloudWatch.

AWS DataSync

DataSync provides status information about each task, task execution, and whether the locations that it is configured for (such as on-premises NFS servers and EFS file systems) are reachable. After you create a task, you can easily tell whether it's available to run, whether it's unavailable to run because one or more of the file systems are unreachable, or whether it's currently running. If the task is running, you can inspect the task execution status to see what phase of the workflow it is in, and ultimately if it succeeded or failed.

If you want to integrate DataSync operational health information into a custom monitoring solution, you can use the DataSync API to obtain a task's status or a task's execution status. You also can enable DataSync to [send task logging information](#) to an Amazon CloudWatch Log Group to monitor and debug your tasks.

DataSync also sends event and health metrics to Amazon CloudWatch. Using CloudWatch Events you can monitor and configure alarms for state changes in an agent, location, task, or task execution. Using CloudWatch metrics, as a function of time, you can see the number of files written, the number of bytes written, and the number of bytes transferred over the network as a function of time. For more information, see [Monitoring Your Task](#).

Amazon EFS

As discussed previously, EFS has two performance modes: General Purpose and Max I/O. General Purpose mode is designed to be the right fit for most applications, but if you require the highest amounts of aggregate throughput possible, you might be better served by Max I/O mode. To determine whether your workload is a good fit for Max I/O, we recommend that you first test using General Purpose mode and monitor the EFS PercentIOLimit CloudWatch metric. If you see this value consistently at or near 100%,

then a Max I/O file system might be a better option for you. Another option would be to split your workload across multiple General Purpose mode file systems.

If you're not reaching the I/O limit of General Purpose mode but still want to drive additional throughput, there are two techniques that you can pursue. First, to take maximum advantage of the distributed design of EFS, where storage is replicated within and across AZs over an unconstrained number of servers, EFS recommends leveraging parallel I/O as much as possible. EFS [provides a tutorial hosted on GitHub](#) with several examples showing how to put this technique into practice, including an AWS CloudFormation stack to help you get up and running.

A second option is to determine if your throughput is being rate limited based on your storage. This might be the case if you're running in EFS Bursting Throughput mode, have a relatively small file system, and have run out of burst credits. Each file system begins with 2.1 TB of burst credits, which accrue when I/O is idle but are consumed as I/O is performed. If this is the case and you have deterministic application throughput requirements, consider EFS Provisioned Throughput.

EFS [provides CloudWatch metrics](#) that allow you to monitor your throughput, burst credits, and client connections.

Pricing and Cost Considerations

AWS DataSync

AWS DataSync has simple, per-GB pricing. You pay \$0.04/GB or \$40/TB transferred in all AWS Regions where DataSync is available.

Amazon EFS

Amazon EFS has four price dimensions: Standard Storage, Infrequent Access storage, Infrequent Access IO requests, and Provisioned Throughput. There are no minimum fees and you pay only for what you use.

Using the US-EAST-1 Region, EFS Standard storage costs \$0.30/GB-month. EFS IA storage costs \$0.025/GB-month, 92% less than EFS Standard. The cost savings of EFS IA is consistent in all Regions where EFS is available.

Using EFS IA, you also pay a fee each time you read a file from the IA storage class, and each time lifecycle management moves a file to the IA storage class. This fee is \$0.01/GB transferred (US-EAST-1).

You also have the option of provisioning throughput for EFS at the rate of \$6/MB/s (US-EAST-1). You only pay for the throughput in excess of what your file system is permitted to drive, based on Standard storage, when in bursting throughput mode. As an example, if you have a file system with 10 TiB in the Standard storage class, it is permitted to drive 500 MB/s of throughput continuously. If you provision 1 GB/s of throughput, you pay only for the additional 500 MB/s relative to what's permitted from your storage.

AWS Backup

Backing up an EFS file system costs \$0.05/GB-month for warm storage and \$0.01/GB-month for cold storage (US-EAST-1). Restores cost \$0.02/GB from warm storage and \$0.03/GB from cold storage.

Migration Cost Example

Here is an example for calculating migration costs. Ultimately, your workload and access patterns will drive your actual cost, so we've made simplifying assumptions based on generally accepted industry estimates. To test your own assumptions and analysis, refer to the [AWS Pricing Calculator](#).

Let's say that you are migrating 100 TB of storage from an on-premises NFS server to AWS and storing it for 1 year, and 80% of this storage is infrequently accessed, aging off and never read again. You also want to keep a daily backup of your file system. Your yearly cost will be:

Component	Cost
AWS DataSync	$\$0.04/\text{GB} * 1000 \text{ GB/TB} * 100 \text{ TB} = \mathbf{\$4,000}$
Amazon EFS Standard Storage	$\$0.30/\text{GB-month} * 1000 \text{ GB/TB} * 100 \text{ TB} * 20\% \text{ frequently accessed} * 12 \text{ months/year} = \mathbf{\$72,000}$
Amazon EFS IA I/O	$\$.01/\text{GB} * 1000 \text{ GB / TB} * 100 \text{ TB} * 80\% \text{ infrequently accessed} = \mathbf{\$800}$
Amazon EFS IA Storage	$\$0.025/\text{GB-month} * 1000 \text{ GB/TB} * 100 \text{ TB} * 80\% \text{ infrequently accessed} * 12 \text{ months/year} = \mathbf{\$24,000}$
AWS Backup	$\$0.05/\text{GB-month} * 1000 \text{ GB/TB} * 100 \text{ TB} * 12 \text{ months/year} = \mathbf{\$60,000}$

Total cost: **\$160,800** or **\$161/TB/year** to migrate to a highly available, highly durable, elastic, fully managed file system—and keep an extra backup to protect against accidental data deletion.

Conclusion

In summary, this whitepaper provides guidance on how to setup, tune, and monitor migrating applications that use file system storage to AWS. Using AWS DataSync to quickly and efficiently transfer files from an on-premises NFS file system to Amazon EFS and AWS Backup for automated data protection and recovery, you can take advantage of all the benefits of AWS cloud services for around \$161/TB/yr.

Contributors

Contributors to this document include:

- Joe Travaglini, Senior Product Manager, Amazon Elastic File System

Document Revisions

Date	Description
May 2019	First publication
September 2019	Updated costs and reference to additional lifecycle policies