



Decoupling data storage from SAS servers using Amazon FSx for Windows File Server

AWS Prescriptive Guidance



AWS Prescriptive Guidance: Decoupling data storage from SAS servers using Amazon FSx for Windows File Server

Copyright © 2026 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

Introduction	1
Overview	1
Example scenario	1
Targeted business outcomes	3
Data file requirements	4
Architecture design	5
SAS requirements	5
AWS architecture design	5
Architecture components	8
Environment separation	8
Infrastructure components	9
Production environment	9
Integration environment	10
Local storage for SAS servers	10
Shared storage infrastructure	12
Production environment	12
Integration and testing environment	12
Performance	13
Back up and file recovery	13
Disaster recovery	13
End-user benefits	14
Best practices	15
Next steps	16
Resources	17
Document history	18

Decoupling data storage from SAS servers using Amazon FSx for Windows File Server

Katja Mueller and Petar Forai, Amazon Web Services (AWS)

June 2023 ([document history](#))

Overview

This guide provides recommendations and best practices for using [Amazon FSx for Windows File Server](#) to decouple your data storage from SAS servers. The decoupling approach is based on a shared storage architecture that can make your SAS deployment on AWS more resilient by splitting your data and SAS project code into separate applications that you can manage separately in the event of failure or downtime. This architecture can also improve the user experience for end users in your organization by making data handling faster and easier. You can even use the decoupling approach to simplify and streamline your SAS deployment on AWS.

We recommend that you use this guide as a supplement to the AWS Whitepaper [Best Practices for Deploying SAS Server on AWS](#). This guide highlights how you can enhance the architecture described in the whitepaper by using FSx for Windows File Server as a more advanced storage alternative to SAS storage.

Example scenario

This guide is based on an example scenario that assumes the following:

- The source environment is on premises and uses SAS 9.4 for Windows.
- The destination environment is AWS.
- The stored data from the SAS server in the on-premises environment consists of Microsoft Excel files.
- End users access the data in the Excel files directly from their machines without using a central SAS server.
- End users are familiar with using Windows File Server for shared storage and want to continue using Windows Files Server on AWS.

- The file structure for input and output must remain in place permanently and must be referenceable to meet compliance rules and regulations.
- The data from the source environment is migrated from [Amazon Simple Storage Service \(Amazon S3\)](#) to FSx for Windows File Server.

The details of your use case may differ from this example scenario, but you can still use the recommendations and best practices in this guide to design an architecture that decouples data storage from your SAS servers.

Targeted business outcomes

This guide can help you and your organization achieve the following business outcomes:

- Understand the data file requirements for FSx for Windows File Server.
- Learn how to design a decoupled architecture that meets infrastructure requirements for both SAS and AWS.
- Understand the architecture components essential to building a decoupled architecture.
- Learn best practices for decoupling data storage from SAS servers.

Data file requirements

Before you deploy your SAS server on AWS, we recommend that you prepare your data files to meet important requirements. We recommend that your data files are:

- Permanently and consistently referenceable by the same path/location
- Securely available
- Easy to share
- Easy to access for end users
- Recoverable through self-service
- Able to integrate with existing Microsoft Active Directory environments

Amazon FSx for Windows File Server can help you meet all these requirements. FSx for Windows File Server provides fully managed shared storage that's built on Windows Server. FSx for Windows File Server has native support for Windows file system features and for the industry-standard Server Message Block (SMB) protocol that's used to access file storage over a network. Your SAS server can map to FSx for Windows File Server by using the SMB protocol.

A key standard feature of FSx for Windows File Server is encryption at rest. This is important if your organization must handle data that contains highly-sensitive information. Data in FSx for Windows File Server is always encrypted at rest, and you can use your own customer managed keys for encryption. For more information, see [Encryption at Rest](#) in the Amazon for FSx Windows User Guide.

Architecture design

This section provides high-level guidance on how to meet the recommended infrastructure requirements from SAS and how to design your AWS architecture for a SAS deployment.

SAS requirements

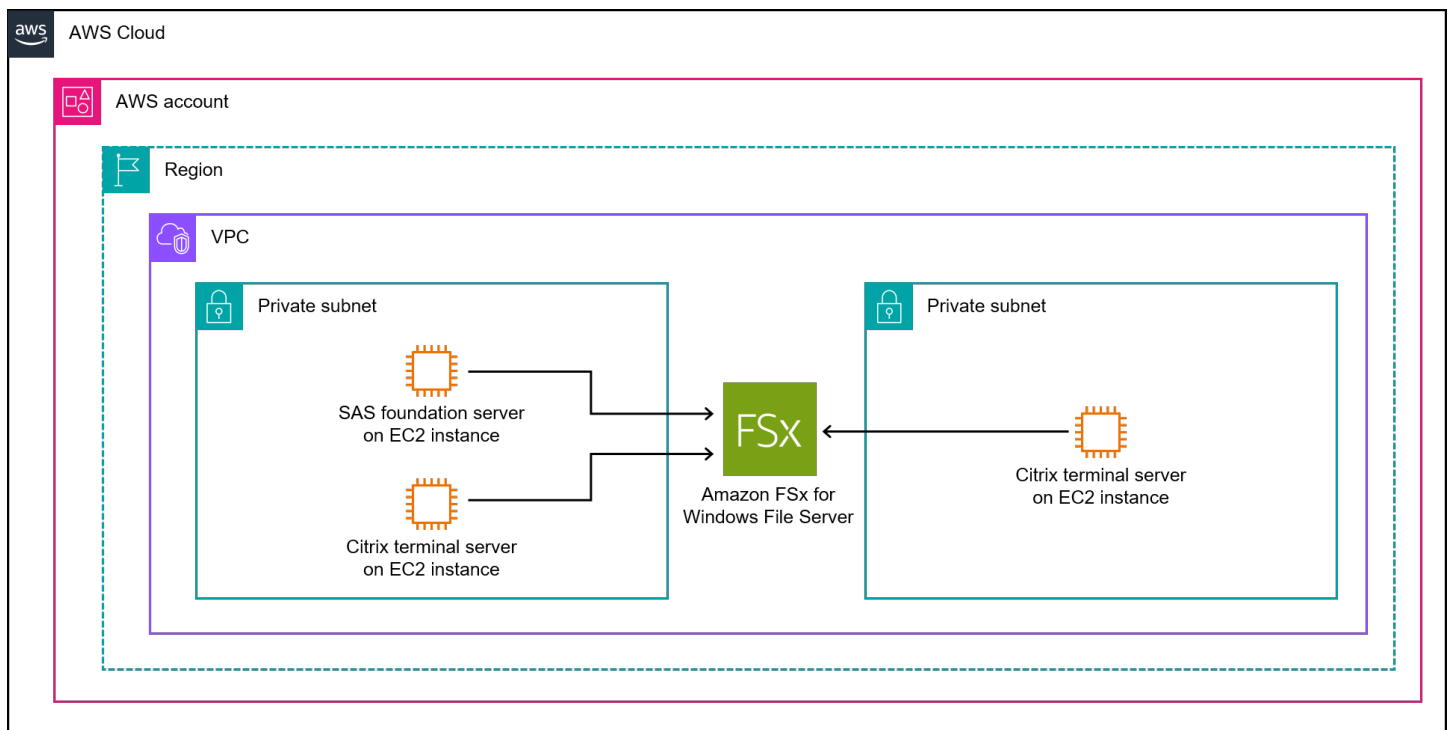
SAS provides official guidelines for all the major technical infrastructure components for deploying SAS 9.4. We recommend that you familiarize yourself with these technical system requirements as part of your migration planning process. For more information regarding SAS I/O requirements, see [Best Practices for Configuring your IO Subsystem for SAS® 9 Applications](#) in the SAS documentation.

AWS architecture design

Designing the right architecture for your SAS environment is critical to the success of your SAS deployment. This section provides an overview of a recommended AWS architecture that separates SAS environments into the following AWS accounts:

- An account dedicated to your SAS production environment
- An account dedicated to your SAS integration environment (for testing)
- A shared account dedicated to infrastructure management

It's a best practice to separate these environments so that you can independently test and validate changes before promoting changes from the test environment to the production environment. The following diagram illustrates the decoupling architecture.



The two accounts include user-facing Citrix terminal servers. End users in your organization can connect to these terminal servers by using a Citrix client or the Citrix web publishing portal, and then launch the SAS Enterprise Guide SAS client. The shared infrastructure management account can include all the automation and CI/CD tooling that's required for deploying your infrastructure components into each AWS account and AWS Region.

The SAS client connects to the SAS server that's running within the same network. Both servers have access to a dedicated high-performance shared storage file server: FSx for Windows File Server. The servers and file server are joined to your on-premises Microsoft Active Directory environment. This enables operational IT to centralize access management. Additionally, you can base the patch management in this environment on your existing on-premises software distribution, patch management solution, and the Citrix terminal server infrastructure.

This architecture doesn't use bastion hosts, implemented as additional Amazon Elastic Compute Cloud (Amazon EC2) instances, for console and direct management. All EC2 instances are configured to use AWS Systems Manager Fleet Manager in case the administrative Citrix access doesn't work. Fleet Manager, a capability of AWS Systems Manager, provides an interface that helps you to remotely manage nodes running on AWS or in on-premises environments. For more information, see [AWS Systems Manager Fleet Manager](#) in the AWS Systems Manager User Guide.

The VPC is part of the landing zone's hub and spoke network. The network provides on-premises, internet breakout, and other external connectivity centrally.


Architecture components

This section outlines the specifications of the following important functional architecture components:

- **SAS server** – This server is the central compute component for analytics processing and includes local direct-attached storage (DAS).
- **SAS subversion server** – This server acts as the centralized version control system for SAS.
- **Amazon FSx for Windows File Server** – This is an SMB file server for sharing storage between the SAS server and terminal servers. End users store and archive their pre- and post-processed data files on FSx for Windows File Server.
- **Microsoft Remote Desktop Services (RDS), also known as Terminal Services** – RDS allows end users to access SAS servers by using a SAS client.
- **Infrastructure automation** – You can use the AWS Cloud Development Kit (AWS CDK) with AWS CodePipeline and AWS CodeCommit to automate your infrastructure. CodePipeline can help you provision your infrastructure components. CodePipeline is a continuous delivery service for modelling, visualizing, and automating the steps required to release code. Additionally, CodePipeline provides a shared central environment and enables infrastructure management that's independent from local machines. CodeCommit is a secure, highly scalable, fully managed source control service that hosts private Git repositories. You can use CodeCommit to store AWS CDK infrastructure automation code and parameters.

Environment separation

The following diagram shows an architecture for separating a SAS integration and SAS production environment.



Infrastructure components

This section provides an overview of the infrastructure components that are required for the recommended architecture in this guide.

Production environment

We recommend that you use the following infrastructure components for your production environment.

Type	Instance type	Resources
1 SAS server	m6i.4xlarge	16 vCPUs (8 cores) 64 GB RAM
2 Citrix terminal servers	m6i.4xlarge	16 vCPUs (8 cores) 64 GB RAM (for example, 1–2 GB per user session for Microsoft Office and Adobe)

Suite and 500–1024 MB per SAS client on average)

25+ users

Potential to scale out with more terminal servers in the future

1 SAS subversion server	m6i.2xlarge	8 vCPUs
		4 cores
		32 GB RAM

Integration environment

We recommend that you use the following infrastructure components for your integration environment.

Type	Instance type	Resources
1 SAS server	m6i.2xlarge	8 vCPUs (4 cores) 32 GB RAM
2 terminal servers	m6i.2xlarge	8 vCPUs (4 cores) 32 GB RAM
1 SAS subversion server	m6i.xlarge	4 vCPUs (2 cores) 16 GB RAM

Local storage for SAS servers

The recommended architecture uses M6i instances based on the latest Intel Xeon Scalable processors and uses the Nitro Hypervisor from the [AWS Nitro System](#). The M6i instance type

is optimized for [Amazon Elastic Block Store \(Amazon EBS\)](#) and offers dedicated bandwidth for network-accessed EBS volumes. The following table includes details about instance storage configuration for non-shared storage. You can attach additional EBS volumes on demand.

Server	Type	Capacity	Production	Testing
SAS server	Storage type	AWS resource/ service and EBS type	Requirement on seq. IO (read/wri te)	Same as production
SAS server	Operating system boot and swap	EBS 200 GB (gp3)	Not relevant for sizing due to low requirements	Same as production
SAS server	SASWORK	EBS 2x 512 GB (gp3/each 5,000 IOPS) in RAID 0	8 * 150 Mbps, 1200 Mbps or ~ 11.5 Gbps M6i instance support 12.5 Gbps EBS storage bandwidth with gp3 EBS volumes	1x 1024 GB volume gp3 5,000 IOPS
SAS server	SAS Software Depot and other auxiliary storage (to include SAS install in addition)	EBS 125 GB (gp3)	Not relevant for sizing due to low requirements	Same as production
SAS terminal server	Operating system boot and swap	EBS 100 GB (gp3)	Not relevant for sizing due to low requirements	Same as production

SAS SVN server	Operating system boot and swap	EBS 100 GB (gp3)	Not relevant for sizing due to low requirements	100 GB
SAS SVN server	Subversion repositories	EBS 1000 GB (gp3)	Default	400 GB in addition to ops drive

Shared storage infrastructure

We recommend using FSx for Windows File Server as a shared storage solution for your SAS server and the Citrix terminal servers. You don't have to use S3 buckets for any additional file storage, unless you need the bucket for maintaining system information or automation scripts.

You can also store the subversion checkout/working copy of project code on FSx for Windows File Server. The SAS subversion server stores the repositories locally. The subversion server acts as the central version control system.

We recommend that you use FSx for Windows File Server to store Windows user profiles across your Citrix terminal servers. This will enable seamless load balancing across both servers.

Production environment

The architecture in this guide is designed to meet the following requirements for the production environment:

- **Storage type** – FSx for Windows File Server
- **Type** – Multiple Availability Zones
- **Resource/throughput** – 1024 MB
- **Storage** – 1.2 TB SSD

Integration and testing environment

The architecture in this guide is designed to meet the following requirements for the integration environment:

- **Storage type** – FSx for Windows File Server

- **Type** – Multiple Availability Zones
- **Resource/throughput** – 512 MB
- **Storage** – 512 GB SSD

Performance

The I/O throughput for FSx for Windows File Server is easy to adjust, and you can build I/O throughput dashboards to meet your monitoring needs. You can also enable the operations team to adjust throughput based on end-user needs.

Back up and file recovery

All SAS data resides on a separate FSx for Windows File Server as persistent storage. There are two levels of backup implemented on the data stored in FSx for Windows File Server:

1. **Daily backups retained for 30 days** – These backups are retained in an S3 bucket. You can use this snapshot-based backup for recovery if an Amazon FSx volume is corrupted or lost.
2. **Backups retained using Microsoft Volume Shadow Copy Service (VSS)** – Files on the FSx for Windows File Server are snapshotted for backup to a special storage partition on the FSx for Windows File Server twice per day and retained indefinitely. The backup is based on available storage of the VSS partition on FSx for Windows File Server (up to 10 percent of total storage space). If end users corrupt or lose a file on FSx for Windows File Server, they can initiate their own restore directly from Windows File Explorer on the SAS terminal servers.

Disaster recovery

The decoupling architecture in this guide is designed with disaster recovery in mind. Amazon FSx is deployed across two AWS Availability Zones. If the Availability Zone where the active FSx for Windows File Server resides become unavailable, then the service automatically fails over and provides the file sharing services from the second Availability Zone.

End-user benefits

An architecture that decouples your data storage from SAS servers can improve the overall speed and usability of data access. For example, end users can access, share, prepare, and handle files significantly faster by using a decoupled architecture instead of a traditional architecture that relies on individual desktop devices. The use case for this guide achieved 35 percent time savings compared to the previous setup. A central terminal server deployment, based on Citrix clients for the scenario used in this guide, offers the following advantages:

- Offers portability and mobility of work processes for each end user
- Enables end users to resume working sessions
- Keeps data secure in the cloud away from local device storage that can be lost
- Offers continuous availability by using staged rollouts of patches and backups

Best practices

We recommend the following best practices when you design your decoupled data architecture:

- **Plan your data migration** – Develop a detailed plan on how to migrate data from a single production environment to a multi-environment setup early on in the design process. This can help you identify bottlenecks, reduce risks, and properly allocate the time and resources that are required for the migration.
- **Keep the design simple and modular** – The example scenario in this guide is based on a multi-account architecture that uses three different AWS accounts. It's typically not necessary to use more than one AWS account per environment. We recommend that you place all infrastructure for the same environment in the same AWS account. Then, you can use the automation capabilities of AWS CDK to operationalize restore processes.
- **Design your architecture to be secure and compliant** – We recommend that you design your architecture to meet the compliance and operational best practices required by your organization. For example, you can design your architecture to meet 21 CFR Part 11 electronic signature requirements.
- **Minimize cost** – Throughput is a significant factor in the cost of a data migration. This is because SAS strongly recommends high I/O throughput from SAS servers to any file servers. We recommend that you review and assess these requirements for your unique use case.
- **Simplify future operations and build resilience** – We recommend that you use FSx for Windows File Server to simplify your operations and build resilience. This is possible because FSx for Windows File Server offers the following benefits:
 - No patching, agents, or antivirus software required
 - No manual failover involved
 - Self-service recovery of deleted or corrupted data files for end users
 - Easier to share data across other potential data analysis applications and systems
 - Simplifies the management of servers (servers hold the application but data resides elsewhere)

Next steps

We recommend that you take the following next steps to decouple your data storage from your SAS servers:

1. Monitor your SAS application to understand usage pattern and data volumes.
2. Understand how Amazon FSx for Windows File Server works and its cost structure.
3. Plan sufficient time for pre-migration testing.
4. Plan for a maintenance window to complete the migration.

Resources

- [Best Practices for Deploying SAS Server on AWS](#) (AWS Whitepaper)
- [Best Practices for Configuring Your I/O Subsystem for SAS®9 Applications](#) (SAS documentation)
- [What is FSx for Windows File Server?](#) (Amazon FSx Windows User Guide)
- [Working with shadow copies](#) (Amazon FSx Windows User Guide)
- [Encryption at Rest](#) (Amazon FSx Windows User Guide)
- [AWS Systems Manager Fleet Manager](#) (AWS Systems Manager User Guide)

Document history

The following table describes significant changes to this guide. If you want to be notified about future updates, you can subscribe to an [RSS feed](#).

Change	Description	Date
Initial publication	—	June 20, 2023