

Service Name

Amazon FSx for Windows File Server

Service Overview

Amazon FSx for Windows File Server is a fully managed Microsoft Windows file server service that provides enterprise-level features, performance, and compatibility for Windows workloads. Built on Windows Server, it delivers shared file storage accessible over the industry-standard Server Message Block (SMB) protocol. The service offers high levels of throughput and IOPS with consistent sub-millisecond latencies, supporting thousands of concurrent compute instances and devices. It integrates seamlessly with Microsoft Active Directory for authentication and access control, while providing administrative features like user quotas, end-user file restore, data deduplication, and File Server Resource Manager capabilities.

Key Use Cases

- **Business applications:** Enterprise Windows applications requiring shared file storage with Active Directory integration
- **Home directories and user shares:** Centralized storage for user profiles and personal folders with quota management
- **High availability SQL Server:** Continuously available file shares for SQL Server Always On Failover Cluster Instances
- **Content management and web serving:** Shared storage for content management systems and web applications
- **Virtual desktop infrastructure (VDI):** File storage for Amazon WorkSpaces and AppStream 2.0 environments
- **Software development:** Build environments and source code repositories with version control integration

Features

- **Microsoft Active Directory Integration:** Native integration with AWS Managed Microsoft AD and self-managed Active Directory for authentication and access control
- **Multi-AZ Deployment:** High availability across multiple Availability Zones with automatic failover and standby file servers
- **Data Deduplication:** Automatic identification and elimination of duplicate data, delivering typical savings of 50-60% for general-purpose file shares
- **File Server Resource Manager:** Advanced file management including quotas, file screening, classification, and storage reporting capabilities

- **Shadow Copies:** Point-in-time snapshots enabling end-user file restore and version recovery
- **Performance Scaling:** Independent scaling of storage capacity, throughput capacity, and SSD IOPS with up to 12 GB/s throughput and 400,000 IOPS
- **Hybrid Connectivity:** On-premises access via AWS Direct Connect, Site-to-Site VPN, and FSx File Gateway with local caching

Value Proposition

Amazon FSx for Windows File Server eliminates the complexity and overhead of managing Windows file server infrastructure while providing enterprise-grade features and performance. Customers benefit from fully managed operations including patching, backups, and monitoring, reducing operational burden by up to 80%:

- The service delivers native Windows compatibility with SMB protocol support, ensuring seamless integration with existing Windows applications without code changes
- Cost optimization features like data deduplication, flexible storage types (SSD/HDD), and pay-as-you-go pricing provide significant savings compared to on-premises solutions
- Built-in security features including encryption at rest and in transit, integration with AWS CloudTrail for auditing, and granular access controls meet enterprise security requirements while maintaining high availability through Multi-AZ deployments with 99.99% SLA

Service Integration

Amazon FSx for Windows File Server integrates deeply with core AWS services to provide comprehensive enterprise file storage solutions. It seamlessly works with Amazon EC2 for compute workloads, Amazon ECS for containerized applications, and VMware Cloud on AWS for hybrid environments:

- Integration with AWS Directory Service provides managed Active Directory capabilities, while AWS Backup offers centralized backup management across multiple AWS services
- The service connects with Amazon VPC for network isolation and security groups for access control
- AWS DataSync enables efficient data migration and replication, while AWS Direct Connect and Site-to-Site VPN facilitate hybrid connectivity
- Integration with Amazon CloudWatch provides monitoring and metrics, AWS CloudTrail enables audit logging, and AWS Identity and Access Management (IAM) controls administrative access

Onboarding and Offboarding

Customers can get started with Amazon FSx for Windows File Server through the AWS Management Console, AWS CLI, or API. The process involves creating an AWS account, setting up Microsoft Active Directory (either AWS Managed or self-managed), launching Amazon EC2 instances, and creating the FSx file system through the console wizard:

- The getting started guide provides step-by-step instructions including joining instances to the directory, configuring security groups, mapping file shares, and testing connectivity
- Prerequisites include an AWS account, VPC configuration, and Active Directory setup
- The service offers tutorials and documentation for common scenarios including SQL Server integration, WorkSpaces deployment, and hybrid connectivity setup
- Initial file system creation typically takes 30-45 minutes with immediate access to administrative features

Getting Started Guide: <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/getting-started.html>

Data Removal

Data removal requires deleting the Amazon FSx file system through the AWS console, CLI, or API. Before deletion, customers should back up important data and update applications to remove file system dependencies. The deletion process permanently removes all data and cannot be undone. Customers can optionally create final backups before deletion, which can be retained independently. Cross-region backup copies provide additional protection. The process includes removing DNS aliases, updating security group references, and cleaning up any associated resources like Directory Service connections.

Backup/Restore and Disaster Recovery

Amazon FSx provides comprehensive backup and disaster recovery through automatic daily backups stored in Amazon S3 with 99.999999999% durability. User-initiated backups can be created anytime and retained indefinitely. AWS Backup integration enables centralized backup management across multiple AWS services. Cross-region and cross-account backup copying supports disaster recovery strategies. Backups are file-system-consistent, incremental, and encrypted. Restoration creates new file systems populated with backup data through lazy-loading for faster recovery times.

Backup Capabilities

Amazon FSx offers native backup capabilities with automatic daily backups enabled by default during a configurable backup window. Backups are incremental, file-system-consistent, and stored durably in Amazon S3:

- Users can create additional manual backups anytime through the console, CLI, or API
- AWS Backup integration provides centralized backup management with advanced features like cross-region copying, lifecycle management, and compliance reporting
- Backup retention can be configured from 0-90 days for automatic backups, while user-initiated backups are retained indefinitely until manually deleted

Disaster Recovery

Amazon FSx supports disaster recovery through Multi-AZ deployments providing automatic failover within regions, typically completing in under 30 seconds. Cross-region disaster recovery is achieved through backup copying and AWS DataSync replication:

- The service integrates with AWS Backup for centralized disaster recovery planning across multiple services
- Multi-AZ file systems maintain synchronous replication between primary and standby file servers
- Cross-region replication using native Microsoft DFS Replication or AWS DataSync enables geographic disaster recovery with customizable recovery point objectives

Recovery Objectives

Amazon FSx helps customers meet RPO and RTO objectives through multiple recovery mechanisms. Multi-AZ deployments provide near-zero RPO with synchronous replication and RTO under 30 seconds for automatic failover. Daily backups enable point-in-time recovery with RPO up to 24 hours. More frequent user-initiated backups can reduce RPO to minutes. Cross-region backup copying extends disaster recovery capabilities. File system restoration from backups typically matches new file system creation time, with lazy-loading reducing initial RTO for data access.

Service Constraints

Service Quotas: <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/limits.html>

FAQ: <https://aws.amazon.com/fsx/windows/faqs/>

Technical Requirements

Service Documentation: <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html>

User Guide: <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/getting-started.html>

API Reference: <https://docs.aws.amazon.com/fsx/latest/APIReference/welcome.html>

Pricing Overview

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Public Pricing: <https://aws.amazon.com/fsx/windows/pricing/>

Cost Optimization

Amazon FSx offers several unique cost optimization features including data deduplication delivering typical savings of 30-80% through automatic duplicate data elimination. Customers can choose between SSD and HDD storage types based on performance requirements, with HDD providing significant cost savings for less performance-sensitive workloads:

- Flexible throughput capacity scaling allows independent adjustment of performance and cost
- Storage capacity can be increased as needed, avoiding over-provisioning
- Single-AZ deployments provide approximately 40% cost savings compared to Multi-AZ for non-critical workloads
- User and folder-level quotas help control storage consumption and costs
- The pay-as-you-go model eliminates upfront hardware investments and licensing fees

Savings Considerations

Cost savings come from eliminating on-premises infrastructure, reducing operational overhead by up to 80%, and leveraging AWS economies of scale. Data deduplication typically reduces storage costs by 50-60% for general-purpose file shares:

- HDD storage costs significantly less than SSD while meeting performance requirements for many workloads
- Single-AZ deployments offer substantial savings for development and non-critical environments
- Cross-region data transfer costs should be considered for disaster recovery scenarios
- AWS Backup provides cost-effective long-term retention compared to traditional backup solutions
- The managed service model reduces staffing costs for Windows file server administration, patching, and maintenance
- Integration with other AWS services can provide additional cost efficiencies through consolidated billing and volume discounts