Service Name

Amazon FSx

Service Overview

Amazon FSx is a fully managed file storage service that makes it easy to launch and run popular file systems without hardware provisioning, software configuration, patching, or backups. FSx provides cost-efficient capacity with high levels of reliability and integrates with other AWS services. It offers four file system types: FSx for Lustre (high-performance for HPC/ML workloads), FSx for Windows File Server (Windows-based SMB file shares), FSx for NetApp ONTAP (enterprise features with NFS/SMB support), and FSx for OpenZFS (Linux workloads with ZFS features). Each file system type is optimized for specific workloads and provides enterprise-grade performance, security, and availability features with flexible deployment options including Single-AZ and Multi-AZ configurations.

Key Use Cases

- High Performance Computing (HPC): FSx for Lustre provides up to 1000 GB/s throughput for machine learning, financial modeling, and scientific computing workloads
- Windows workloads: FSx for Windows File Server supports business applications, home directories, SQL Server deployments, SharePoint, and .NET applications with Active Directory integration
- Content management and media processing: All FSx types support content repositories, video processing, and media workflows with high-throughput file access
- Data analytics and machine learning: FSx for Lustre and ONTAP provide high-performance storage for big data analytics, ML training, and data lake architectures
- Enterprise file shares and collaboration: FSx for Windows File Server and ONTAP provide secure, managed file shares with user quotas and access controls
- Backup and archival: Integration with AWS Backup and S3 for automated backup strategies and long-term data retention

Features

- **High Performance Storage**: Sub-millisecond latencies, up to 1000 GB/s throughput for Lustre, hundreds of thousands of IOPS, and scalable performance independent of storage size
- **Multi-Protocol Support**: SMB 2.0-3.1.1, NFS 3.0-4.2, and iSCSI protocols with cross-platform compatibility for Windows, Linux, and macOS clients
- **Data Management**: Data deduplication, compression, intelligent tiering, user/group quotas, snapshots, and automated lifecycle management

- **High Availability**: Multi-AZ deployments with 99.99% SLA, automatic failover, and disaster recovery capabilities across regions
- **Security and Compliance**: Encryption at rest and in transit, Active Directory integration, VPC isolation, file access auditing, and compliance certifications
- **AWS Integration**: Native integration with S3, EC2, ECS, EKS, AWS Backup, DataSync, and other AWS services for hybrid workflows

Value Proposition

Amazon FSx eliminates the complexity and overhead of managing file storage infrastructure while providing enterprise-grade performance and features. Customers benefit from fully managed operations with no hardware provisioning, automatic software updates, and built-in high availability:

- FSx offers significant cost savings through pay-as-you-go pricing, data deduplication (30-80% savings), and intelligent tiering for cold data
- The service provides superior performance with sub-millisecond latencies and massive throughput scaling, while maintaining compatibility with existing applications and protocols
- Deep AWS integration enables seamless hybrid architectures with S3, automated backups, and native support for compute services like EC2 and containers

Service Integration

Amazon FSx integrates natively with core AWS compute services including Amazon EC2, Amazon ECS, and Amazon EKS for file storage access. AWS Backup provides centralized backup management with automated scheduling and retention policies:

- Amazon S3 integration enables data import/export and tiering for cost optimization
- AWS DataSync facilitates migration and synchronization between on-premises and FSx file systems
- AWS Directory Service and Active Directory provide authentication and access control
- VPC networking enables secure access control with security groups and NACLs
- Additional integrations include AWS CloudTrail for auditing, CloudWatch for monitoring, AWS Direct Connect and VPN for hybrid connectivity, and AWS Lambda for automated workflows

Onboarding and Offboarding

Customers begin by setting up prerequisites including an AWS account, VPC configuration, and Active Directory (for Windows File Server). The getting started process involves creating a file system through the AWS Console, CLI, or API by selecting the appropriate file system type, storage capacity, throughput levels, and network configuration:

- After creation, customers mount the file system to compute instances using standard protocols (SMB, NFS) and begin accessing shared storage
- The service provides step-by-step tutorials and documentation for each file system type
- Initial setup typically takes minutes, and customers can immediately begin using enterprise features like snapshots, backups, and user quotas without additional configuration

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

Data Removal

Customers can delete file systems and all associated data through the AWS Console, CLI, or API. Before deletion, final backups are automatically created and retained according to backup retention policies. All data is permanently removed from FSx storage, and customers are no longer charged for storage or throughput capacity. Existing backups remain available until their retention period expires or they are manually deleted. For complete offboarding, customers should also remove associated resources like mount targets, security groups, and any automated backup policies to ensure no residual costs.

Backup/Restore and Disaster Recovery

Amazon FSx provides comprehensive backup and disaster recovery capabilities including automatic daily backups with configurable retention periods (0-90 days), user-initiated backups retained indefinitely, and integration with AWS Backup for centralized management. All backups are incremental, encrypted, and stored across multiple Availability zones. Cross-region backup copying supports disaster recovery strategies, and point-in-time restore creates new file systems from any backup. FSx also supports native features like Windows shadow copies and NetApp SnapMirror for additional data protection and replication capabilities.

Backup Capabilities

Amazon FSx includes native backup functionality with automatic daily backups enabled by default and user-initiated backups available on-demand. All backups are incremental, file-system-consistent, and encrypted:

- FSx integrates seamlessly with AWS Backup for centralized backup management, automated scheduling, and cross-service backup strategies
- AWS Backup provides additional features like unlimited retention, cross-region backup copying, and compliance reporting
- Both native FSx backups and AWS Backup support point-in-time recovery and restore to new file systems

Disaster Recovery

Amazon FSx supports comprehensive disaster recovery through Multi-AZ deployments with automatic failover, cross-region backup copying, and replication capabilities. While FSx doesn't directly integrate with AWS Elastic Disaster Recovery, customers can implement DR strategies using FSx backup and restore capabilities combined with infrastructure automation. FSx for ONTAP supports NetApp SnapMirror for real-time replication, and customers can use AWS DataSync for scheduled replication between file systems across regions for robust disaster recovery architectures.

Recovery Objectives

Amazon FSx helps customers meet aggressive RPO and RTO objectives through multiple mechanisms. Multi-AZ deployments provide near-zero RTO with automatic failover in minutes and RPO measured in seconds through synchronous replication. Frequent automated backups enable RPO of hours with RTO dependent on restore times. Cross-region backup strategies support longer-term disaster recovery scenarios. FSx for ONTAP's SnapMirror replication can achieve RPO of minutes with RTO based on failover procedures, while intelligent tiering and local caching minimize data loss exposure.

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/

User Guide: https://docs.aws.amazon.com/fsx/latest/WindowsGuide/

API Reference: https://docs.aws.amazon.com/fsx/latest/APIReference/

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/pricing/

Cost Optimization

Amazon FSx offers several unique cost optimization features including data deduplication that typically saves 30-80% of storage costs by eliminating duplicate data blocks. Intelligent tiering automatically moves infrequently accessed data to lower-cost storage tiers without performance impact:

- HDD storage options provide cost-effective capacity for workloads that don't require SSD performance
- Single-AZ deployments offer approximately 40% cost savings compared to Multi-AZ for non-critical workloads
- User and group quotas prevent storage overuse, while flexible throughput and IOPS provisioning allows customers to pay only for the performance they need
- The service eliminates traditional storage management costs including hardware, maintenance, and administrative overhead

Savings Considerations

Primary cost savings come from eliminating capital expenditure on storage hardware and reducing operational overhead through fully managed operations. Data deduplication provides immediate 30-80% storage cost reduction for typical workloads:

- Intelligent tiering and HDD storage options reduce ongoing costs for cold data
- Pay-as-you-go pricing eliminates overprovisioning waste, while independent scaling of storage and throughput prevents paying for unused performance
- Single-AZ deployments significantly reduce costs for development and non-critical workloads
- Integration with AWS Backup reduces backup infrastructure costs, and efficient data transfer within AWS regions minimizes bandwidth charges
- The elimination of traditional storage maintenance, patching, and management reduces total cost of ownership substantially