

Service Name

Amazon FSx for OpenZFS

Service Overview

Amazon FSx for OpenZFS is a fully managed file storage service that enables users to move data from on-premises ZFS or other Linux-based file servers to AWS without changing application code or data management methods. Built on the open-source OpenZFS file system, it offers highly reliable, scalable, and feature-rich file storage with support for Linux, Windows, and macOS compute instances and containers via the NFS protocol. The service delivers up to 2 million IOPS and 21 GBps of throughput while providing powerful OpenZFS data management capabilities including data compression, snapshots, cloning, multiple volumes, thin provisioning, and user/group quotas.

Key Use Cases

- **Media and Entertainment:** Supporting high-performance workloads for video editing, rendering, and content distribution with low-latency access to large datasets
- **Scientific Research and Analytics:** Enabling genomics, financial data analytics, seismic imagery analysis, and machine learning workloads with high-throughput data processing capabilities
- **Software Development and DevOps:** Providing containerized applications on Amazon EKS with shared storage for development environments, testing, and CI/CD pipelines with instant snapshots and clones
- **Database Applications:** Supporting high-performance database workloads like Oracle with customizable record sizes and compression for optimized storage performance
- **Enterprise File Sharing:** Migrating on-premises ZFS or Linux-based file servers to AWS while maintaining familiar OpenZFS features and management capabilities

Features

- **OpenZFS Data Management:** Transparent compression, continuous integrity verification, snapshots, thin provisioning, and copy-on-write capabilities
- **High Performance:** Up to 2 million IOPS and 21 GBps throughput with sub-millisecond latencies for demanding workloads
- **Multiple Deployment Options:** Multi-AZ (HA), Single-AZ (HA), and Single-AZ (non-HA) deployments with different availability and durability levels
- **Intelligent-Tiering Storage:** Automatically moves data between Frequent Access, Infrequent Access, and Archive tiers based on access patterns with up to 85% cost savings

- **Multi-Protocol Access:** Supports NFS protocol access from Linux, Windows, and macOS clients, with S3 Access Points for object-style access
- **Volume Management:** Multiple volumes per file system with independent storage capacity, compression, NFS exports, and user/group quotas
- **Data Protection:** Automatic daily backups, point-in-time snapshots, on-demand replication across regions/accounts, and AWS Backup integration
- **Security and Compliance:** Encryption at rest and in transit, VPC isolation, IAM integration, and compliance with ISO, PCI DSS, SOC, and HIPAA

Value Proposition

Amazon FSx for OpenZFS provides a compelling value proposition for organizations migrating from on-premises ZFS environments or seeking high-performance NAS storage in the cloud. Unlike traditional file storage solutions, it eliminates the operational overhead of managing file servers, storage volumes, and software patching while preserving familiar OpenZFS features:

- The service offers superior cost optimization through Intelligent-Tiering (up to 85% savings compared to SSD storage), elastic scaling without upfront provisioning, and pay-as-you-go pricing
- With enterprise-grade features like instant snapshots, zero-copy clones, on-demand replication, and sub-millisecond latencies, it delivers both performance and reliability that surpasses many on-premises deployments while providing seamless AWS integration

Service Integration

Amazon FSx for OpenZFS integrates deeply with core AWS services to provide comprehensive cloud solutions. It works natively with Amazon EC2 for compute workloads, Amazon EKS and ECS for containerized applications via CSI drivers, and Amazon S3 through S3 Access Points for object-style data access:

- The service integrates with AWS Backup for centralized backup management across AWS resources, AWS CloudTrail for API logging and compliance, and Amazon CloudWatch for performance monitoring and alerting
- Additional integrations include AWS DataSync for data migration, AWS RAM for cross-account resource sharing, AWS KMS for encryption key management, Amazon VPC for network isolation, and IAM for access control
- These integrations enable customers to build comprehensive, secure, and scalable architectures

Onboarding and Offboarding

Getting started with Amazon FSx for OpenZFS is straightforward through multiple pathways. Customers can use the AWS Management Console's Quick create option for rapid deployment with default configurations, or Standard create for customized setups:

- The process involves setting up an AWS account, creating a file system in their VPC with appropriate security groups and subnets, and mounting the file system from EC2 instances using standard NFS protocols
- AWS provides comprehensive documentation, tutorials, and step-by-step guides
- For migration, customers can use AWS DataSync to transfer existing data from on-premises ZFS systems
- The service supports both programmatic access via AWS CLI/APIs and infrastructure-as-code deployment through CloudFormation templates for automated provisioning

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

Data Removal

When offboarding from Amazon FSx for OpenZFS, customers must first terminate EC2 instances accessing the file system, then delete the file system through the AWS Console, CLI, or API. All associated backups (both automatic and user-initiated) must be explicitly deleted, as they persist independently of the file system. For data migration before deletion, customers can use AWS DataSync, standard file copy operations, or restore from backups to alternative storage. Once deleted, file system data cannot be recovered unless backups exist, emphasizing the importance of data migration planning during offboarding.

Backup/Restore and Disaster Recovery

Amazon FSx for OpenZFS provides comprehensive backup and disaster recovery capabilities through multiple mechanisms. Built-in automatic daily backups are retained up to 90 days with incremental, crash-consistent snapshots. User-initiated backups can be retained indefinitely and copied across regions for disaster recovery. Point-in-time volume snapshots enable file-level recovery and instant clone creation. On-demand data replication supports cross-region and cross-account data synchronization for disaster recovery scenarios. The service integrates with AWS Backup for centralized backup management and supports AWS Backup's logically air-gapped vaults for enhanced data protection.

Backup Capabilities

Amazon FSx for OpenZFS includes robust native backup capabilities with automatic daily backups during configurable backup windows, retained for up to 90 days. Users can create

on-demand backups retained indefinitely, with both backup types being incremental and crash-consistent:

- The service fully integrates with AWS Backup, allowing centralized backup policy management across multiple AWS services
- Backups can be copied across regions for disaster recovery, and customers can restore entire file systems from any available backup
- Additionally, point-in-time snapshots provide granular recovery options at the volume level

Disaster Recovery

Amazon FSx for OpenZFS supports comprehensive disaster recovery through multiple deployment options and data replication capabilities. Multi-AZ deployments provide automatic failover within 60 seconds across Availability Zones:

- On-demand data replication enables cross-region and cross-account disaster recovery with snapshot-based data transfer
- The service supports backup copying across regions and integrates with AWS Backup for disaster recovery orchestration
- While not directly integrated with AWS Elastic Disaster Recovery, customers can implement disaster recovery strategies using backup restoration, on-demand replication, and Multi-AZ deployments for business continuity

Recovery Objectives

Amazon FSx for OpenZFS helps customers achieve aggressive RPO and RTO objectives through multiple mechanisms. Multi-AZ deployments provide RTO under 60 seconds with synchronous replication for near-zero RPO. Point-in-time snapshots enable RPO measured in minutes to hours depending on snapshot frequency. Cross-region backup copying and on-demand replication support disaster recovery scenarios with configurable RPO/RTO based on replication frequency. Automatic daily backups provide baseline recovery capabilities, while instant snapshots and clones enable rapid recovery for testing and development scenarios.

Service Constraints

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

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

Technical Requirements

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

User Guide: <https://docs.aws.amazon.com/fsx/latest/OpenZFSGuide/>

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/openzfs/pricing/>

Cost Optimization

Amazon FSx for OpenZFS offers unique cost optimization features centered around Intelligent-Tiering storage, which automatically moves data between three access tiers (Frequent, Infrequent, Archive) based on usage patterns, delivering up to 85% cost savings compared to SSD storage. The service eliminates upfront hardware investments and provides pay-as-you-go pricing with no minimum commitments:

- Built-in compression reduces storage consumption, while elastic storage scaling with Intelligent-Tiering means customers only pay for data stored, not provisioned capacity
- Optional SSD read cache provides performance optimization without full SSD storage costs
- Flexible throughput capacity allows right-sizing performance to workload requirements, and efficient snapshot and cloning capabilities minimize storage overhead for development and testing environments

Savings Considerations

Primary cost savings come from eliminating on-premises hardware refresh cycles, reducing operational overhead through fully managed service delivery, and leveraging Intelligent-Tiering's automatic cost optimization. Customers save up to 20% compared to traditional on-premises HDD deployments and dramatically more versus SSD storage through intelligent data tiering:

- The service eliminates costs associated with maintaining file server infrastructure, software licensing, and dedicated storage administration
- Efficient data compression, thin provisioning, and instant snapshots reduce storage consumption
- Cross-region replication and backup capabilities provide disaster recovery at lower costs than traditional solutions
- The elastic scaling model ensures customers avoid over-provisioning costs while maintaining performance during peak demand periods