Service Name

Amazon FSx for Lustre

Service Overview

Amazon FSx for Lustre is a fully managed, high-performance Lustre file system service designed for applications requiring fast storage and low latency. It provides sub-millisecond latencies, up to multiple terabytes per second of throughput, and up to millions of IOPS. The service eliminates the traditional complexity of setting up and managing Lustre file systems, enabling users to spin up battle-tested high-performance file systems in minutes. FSx for Lustre is POSIX-compliant and works with existing Linux-based applications without modifications, supporting workloads where storage speed matters most.

Key Use Cases

- Machine learning and AI workloads: Training large-scale models requiring high-throughput data processing and low-latency storage access
- High Performance Computing (HPC): Scientific computing, weather forecasting, seismic imaging, genomic analysis, and financial modeling requiring parallel file system performance
- Media and entertainment: Video processing, rendering, and content creation workflows requiring fast file I/O and high bandwidth
- Electronic Design Automation (EDA): Semiconductor design and verification workloads with high metadata operations on numerous small files
- Big data analytics: Processing massive datasets requiring scalable, high-speed file system interface with integration to long-term data repositories

Features

- **Multiple Deployment Options**: Scratch file systems for temporary storage and shorter-term processing, Persistent file systems for longer-term storage with data replication and automatic server replacement
- **Storage Classes**: SSD for high IOPS workloads, HDD for throughput-intensive workloads, Intelligent-Tiering for automatic cost optimization with frequent, infrequent, and archive access tiers
- **Amazon S3 Integration**: Native integration with S3 buckets for seamless data import/export, automatic synchronization, and data repository associations
- **Scalable Metadata Performance**: User-provisioned metadata IOPS up to 15x higher than automatic allocation, independent of storage capacity for workloads with many small files

- **Elastic Fabric Adapter (EFA) Support**: Enhanced network performance for MPI applications and tightly coupled HPC workloads with high-performance networking
- **Data Compression**: LZ4 compression support to reduce storage consumption and optimize costs while maintaining performance

Value Proposition

Amazon FSx for Lustre provides the only fully managed Lustre file system in the cloud, delivering enterprise-grade performance without operational overhead. It offers up to 34% better price-performance than on-premises HDD storage and up to 70% better than other cloud-based Lustre solutions:

- The service eliminates complex setup and management tasks while providing native S3 integration for seamless data workflows
- With multiple deployment options and storage classes, customers can optimize for both performance and cost
- The Intelligent-Tiering storage class provides fully elastic scaling and automatic cost optimization, making it ideal for workloads with varying access patterns

Service Integration

Amazon FSx for Lustre integrates natively with Amazon S3 for data repository access, enabling automatic import/export and synchronization of datasets. It works seamlessly with Amazon EC2 instances, Amazon ECS Docker containers, and Amazon EKS containers for compute access:

- The service integrates with AWS services like Amazon SageMaker AI for machine learning workflows, AWS Batch for batch processing jobs, and AWS ParallelCluster for HPC workloads
- It also supports AWS Backup for centralized backup management, AWS KMS for encryption key management, and Amazon CloudWatch for monitoring and logging
- On-premises integration is available through AWS Direct Connect and VPN connections

Onboarding and Offboarding

Customers can get started by creating an FSx for Lustre file system through the AWS Console, CLI, or API after setting up an AWS account and configuring necessary IAM permissions. The process involves selecting deployment type (scratch or persistent), storage class (SSD, HDD, or Intelligent-Tiering), storage capacity, and throughput requirements:

- Optional configurations include linking to S3 data repositories, provisioning metadata IOPS, and enabling SSD read cache

- Once created, the file system can be mounted on Amazon EC2 instances using standard Lustre client tools
- The service includes comprehensive documentation, getting started guides, and sample applications to accelerate deployment

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

Data Removal

When offboarding, customers can delete FSx for Lustre file systems through the AWS Console, CLI, or API, which permanently removes all data stored on the file system. Before deletion, data can be exported to linked S3 data repositories or backed up using AWS Backup. For persistent file systems, automatic and user-initiated backups provide data protection options. Customers should ensure all mounted clients are unmounted before deletion and consider copying critical data to alternative storage locations.

Backup/Restore and Disaster Recovery

FSx for Lustre supports automatic daily backups and user-initiated backups for persistent file systems not linked to S3 data repositories. Backups are incremental, file-system-consistent, and highly durable with configurable retention periods from 0-90 days. AWS Backup integration provides centralized backup management with cross-region and cross-account backup copying capabilities. Backup restoration creates new file systems with identical configurations, enabling point-in-time recovery for disaster scenarios.

Backup Capabilities

FSx for Lustre provides comprehensive backup capabilities through automatic daily backups during configurable backup windows and on-demand user-initiated backups. Integration with AWS Backup enables centralized backup management, automated scheduling, unlimited retention options, and cross-region/cross-account backup copying. Backups are incremental to minimize storage costs and support restoration to new file systems with identical deployment types and configurations.

Disaster Recovery

Disaster recovery is supported through backup and restore capabilities, cross-region backup copying, and S3 data repository integration. AWS Backup enables cross-region disaster recovery with logically air-gapped vaults for enhanced protection. The service supports restoration from backups to create new file systems in different regions, and S3 integration provides additional data durability and geographic distribution for critical datasets.

Recovery Objectives

FSx for Lustre helps customers achieve RPO objectives through automatic daily backups with configurable retention and real-time S3 synchronization for linked data repositories. RTO objectives are met through rapid backup restoration capabilities, where new file systems can be created from backups with data lazily loaded for faster availability. Cross-region backup copying and S3 integration provide additional recovery options for meeting stringent business continuity requirements.

Service Constraints

Service Quotas: https://docs.aws.amazon.com/fsx/latest/LustreGuide/limits.html

FAQ: https://aws.amazon.com/fsx/lustre/faqs/

Technical Requirements

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

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

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/lustre/pricing/

Cost Optimization

FSx for Lustre offers several unique cost optimization features including Intelligent-Tiering storage class that automatically moves infrequently accessed data to lower-cost tiers, starting at less than $0.005 per GB-month. The service provides scratch file systems for cost-effective temporary storage, HDD storage options for throughput-intensive workloads at lower costs, and data compression with LZ4 to reduce storage consumption:

- Automatic scaling eliminates expensive overprovisioning, and integration with S3 enables cost-effective long-term data archiving
- Customers can optimize metadata IOPS provisioning and choose appropriate throughput tiers based on workload requirements

Savings Considerations

Main cost savings come from eliminating on-premises infrastructure management overhead and capital expenses while providing up to 34% better price-performance than

traditional HDD storage. The pay-as-you-go model with no upfront costs or minimum commitments reduces financial risk:

- Intelligent-Tiering automatically optimizes storage costs based on access patterns, and incremental backups minimize backup storage expenses
- S3 integration enables customers to store active data on high-performance FSx while archiving cold data in cost-effective S3 storage classes
- Scratch file systems provide significant cost savings for temporary workloads compared to persistent storage alternatives