Amazon FSx for Lustre

High-Performance Cloud Storage for HPC and AI

Zanhua Huang (zanhua@amazon.com)
Duncan Vogel (fvogdunc@amazon.com)

Amazon Web Services



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. Amazon Confidential and Trademark.

What is Amazon FSx for Lustre?

- Fully managed Lustre file system in the cloud
  - No servers to provision, no Lustre expertise required — spin up a production file system in minutes

- Performant
  - Sub-millisecond latencies,
  - Up to terabytes/s of throughput
  - Large metadata IOPS

- Seamless integration with other AWS services: EC2, S3, EKS, ParallelCluster, etc

- Ideal for: AI/ML, Weather Forecasting, Simulations, etc

aws
© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. Amazon Confidential and Trademark.

EFA-Enabled File Systems

Elastic Fabric Adapter (EFA)

- A network Interface launched Nov 2024
- RDMA; Libfabric, NCCL, NIXL, MPI
- Purpose-built for the most powerful GPU instances for AI workload: **P5** (H100), **P6** (B200)
  - Up to 1,200 Gbps throughput per client — 12× more than previous generation
- NVIDIA GPUDirect Storage (GDS) — bypass CPU, move data directly between GPU memory and NIC

Metadata IOPS: up to 192,000 IOPS

Throughput: up to 1000 MB/s per TiB of filesystem storage
Filesystem storage size: PBs

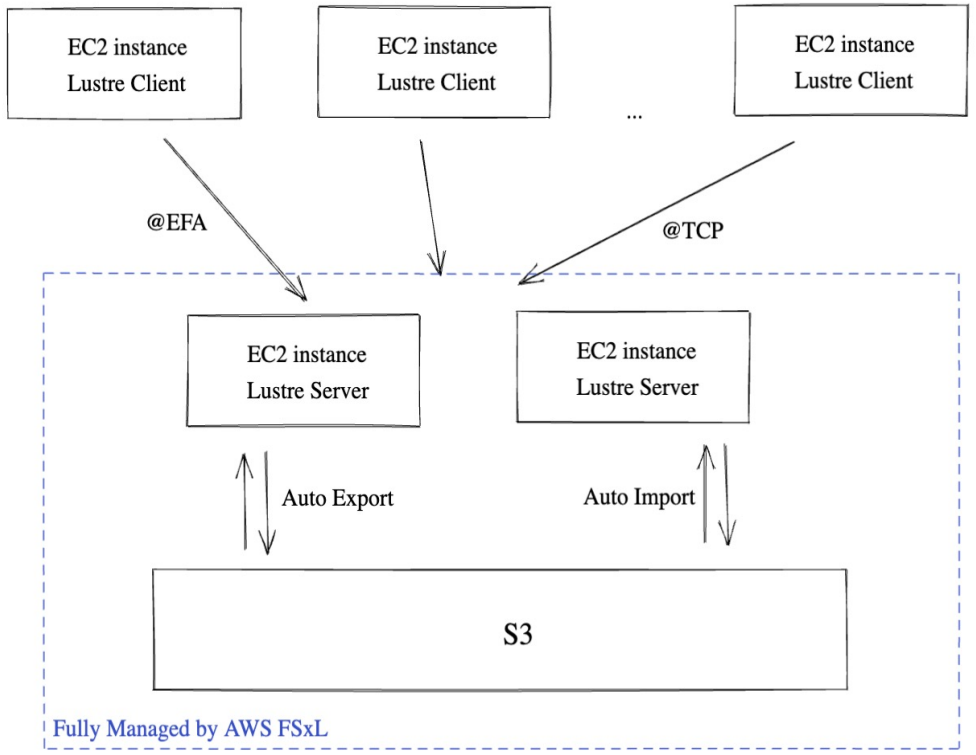
For AI training and HPC workloads
— checkpointing, data loading — run at full network line rate.

aws
© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. Amazon Confidential and Trademark.

HSM with Amazon S3

- Hierarchical Storage Management
- **Auto-export** — auto sync Lustre changes to S3
- **Auto-import** — auto sync S3 changes to Lustre
- Store and access files on S3
- **Ephemeral Lustre**, Durable S3
  - Create — Spin up Lustre linked to S3
  - Compute — Run workloads at Lustre speed
  - Sync — Export changes to S3
  - Tear down — Delete Lustre; data safe in S3
  - Re-create — New Lustre, same data from S3 data

Fully managed by AWS FSx Lustre!



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. Amazon Confidential and Trademark.

FSx Lustre Intelligent-Tiering -- New!

- Available May 2025
- Lustre over ZFS over an Intelligent-Tiering Storage Class
- Intelligent-Tiering**:** automatically moves objects between access tiers based on how frequently they're accessed, without performance impact or operational overhead.
- Performance:
  - Caching— frequently accessed data cached on local SSD
  - 4000 MB/s per OSS;
  - 1 OST/OSS, up to 500 OSSes; up to 500TiB per OSS but pay as you go.
- Pay as you go
  - Storage usage: equals actual file size, no pre-provisioned storage

aws
© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. Amazon Confidential and Trademark.

Open Source Contributions

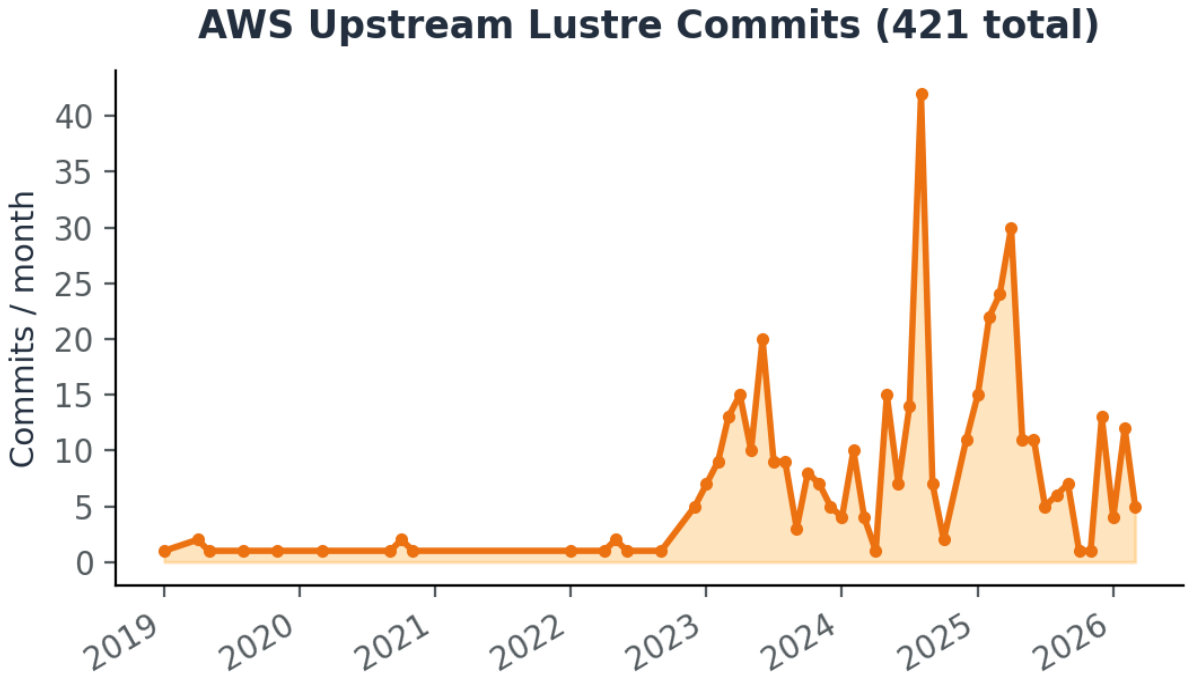
LUG OpenSFS Platinum sponsor

Active contributor to Lustre community

Technical talk at 4:45 PM today:

"Empowering HPC and AI Workloads with optimized Lustre reliability and performance"

— Duncan Vogel & Zanhua Huang




```
git log --pretty=format:"%ae %ad" --date=format:"%Y-%m"
| grep "@amazon.com" | awk '{print $NF}' | sort | uniq -c
```

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. Amazon Confidential and Trademark.

What's Next

- Continued investment in durability, availability, performance, and cost optimization
- Deeper integration with AI/ML services
- Growing our open-source contributions

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. Amazon Confidential and Trademark.

Thank You

Try it: https://aws.amazon.com/fsx/lustre

Deep dive: Join our full talk at 4:45 PM today

"Empowering HPC and AI Workloads with optimized Lustre reliability and performance"

Feedback: We'd love to hear what features matter most to you

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. Amazon Confidential and Trademark.