INNOVATE

FOR EVERY APPLICATION EDITION

25 August, 2022

Next-Generation File Systems:
High performance file-services built for cloud

Gaurav Singh
ISV Solutions Architect
AISPL

Today’s agenda

- Data growth and on-premise challenges
- Overview of Amazon FSx Family
- Demo
- Summary and next steps

Exponential data growth is the new normal

3x

Growth of enterprise data stored by organizations by 2025

87%

of enterprise data will be stored in cloud environment by 2025

>50%

unable to realize full potential value from data

Source: IDC Market Spotlight, Sponsored by AWS, "Organizations Rely on Cloud Storage to Optimize Cost, Increase Agility, and Drive Innovation ," Doc. #USUS48291421, October 2021

aws

Data is used in aspects of your life every day

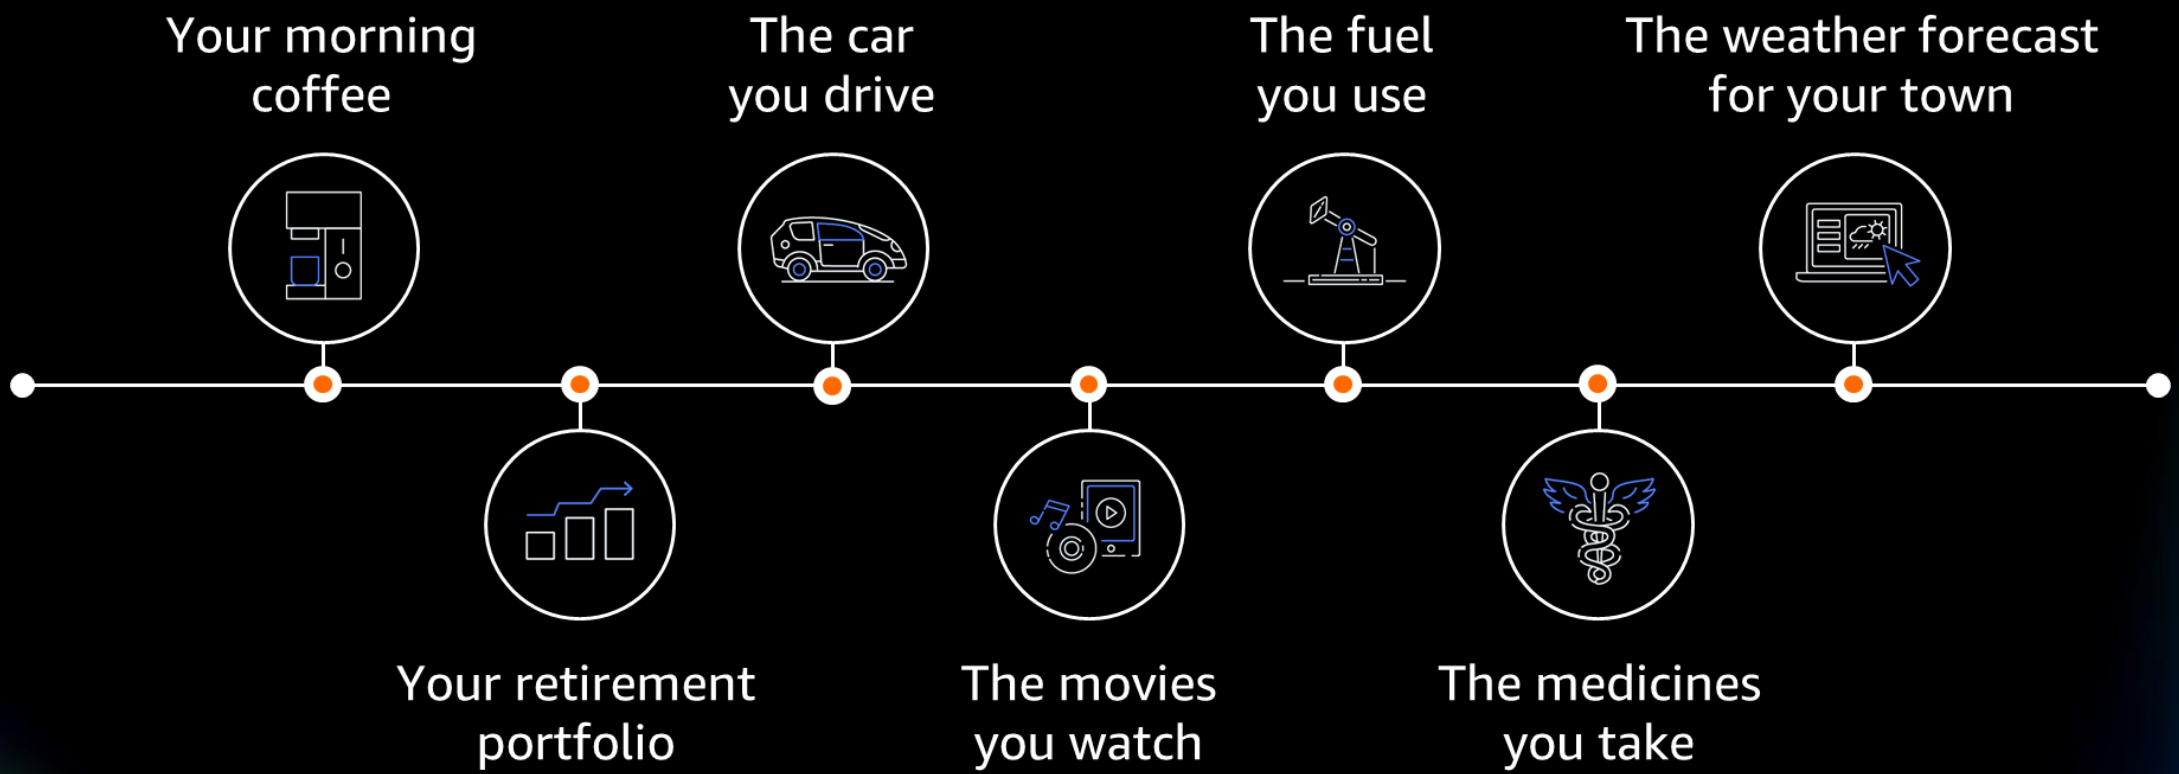

File data is everywhere and growing

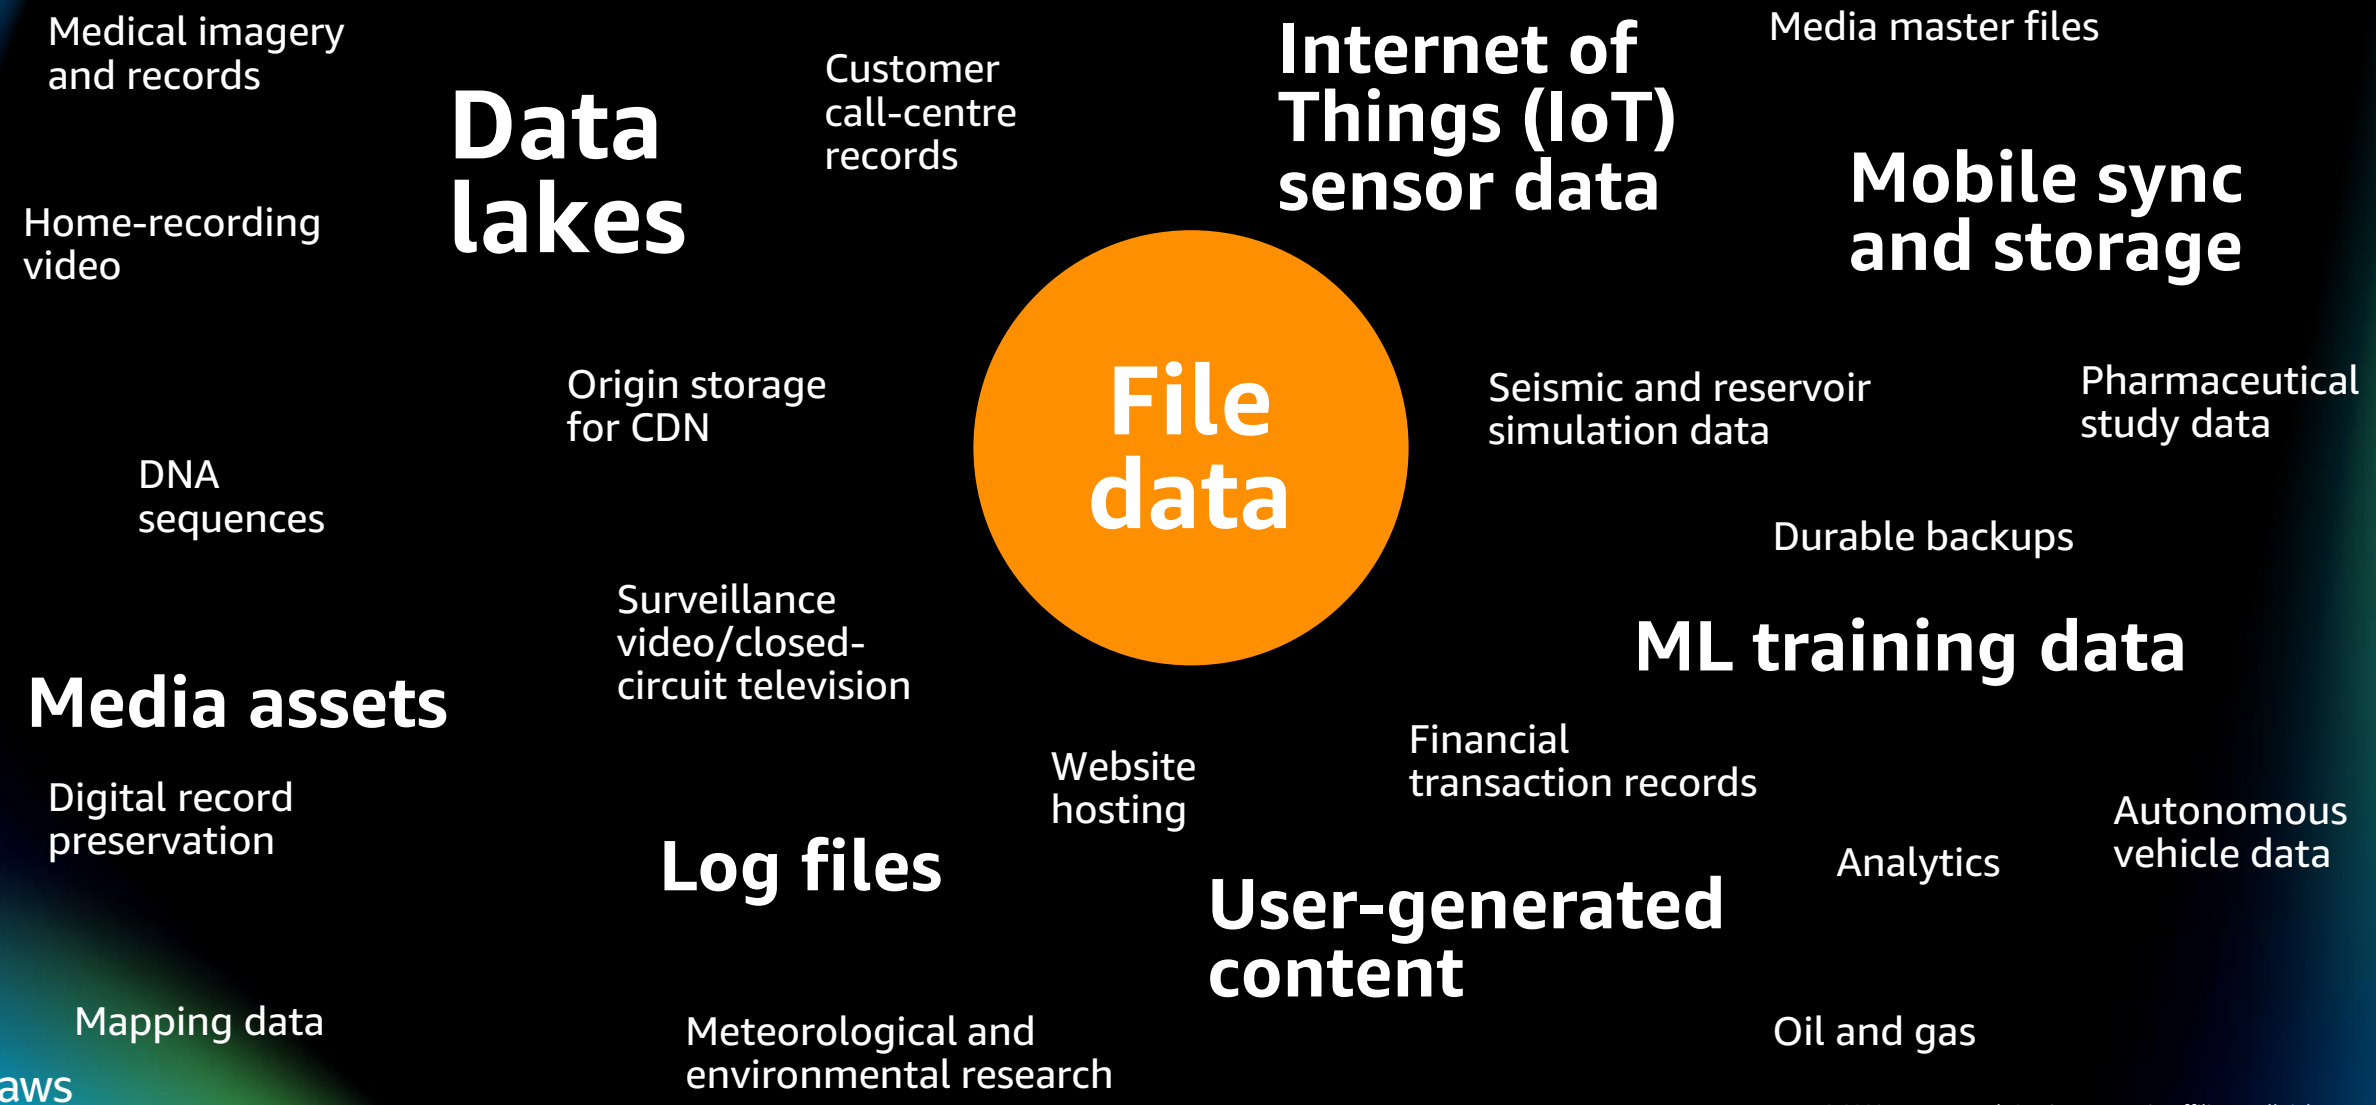

Exponential data growth is a. . .

Challenge

Creating data silos, exceeding ability to scale, and driving higher costs for data

Risk

Caused by increased complexity in protecting data and continuity

Opportunity

To innovate by shifting from **collecting** to **using** data

On-premises challenges

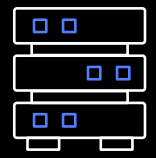

Managing hardware

- Planning capacity
- Periodically procuring and purchasing hardware
- Setting up storage servers and volumes
- Detecting and addressing hardware failures
- High acquisition costs

Managing software

- Installing and configuring server software
- Setting up and configuring file systems
- Applying OS updates
- Managing software licenses
- Managing backups
- Monitoring security

AWS Storage – NextGen infrastructure ready now

FOUNDATIONAL BASIS FOR EVERY APP

Data services

AWS Backup

Amazon EBS snapshots

AWS Transfer Family

Intelligent-Tiering

Compression and dedupe

Lifecycle management

Replication

Metrics reporting

Core services

Object

Amazon S3 and Amazon S3 Glacier

Block

Amazon EBS

File

Amazon FSx Family

Amazon EFS

Hybrid/edge storage

Data movement services

AWS DataSync

AWS Storage Gateway

AWS Snow Family

AWS Outposts

No one-size-fits-all storage solution

Each file system offers customers a unique combination of:

Feature sets

Performance characteristics

Data management capabilities

Cost profiles

Migrating existing workloads to a different storage solution often requires **re-architecting, re-certifying, and re-training**

Customers building new workloads are looking for a **storage solution that's optimized** for their workload

FSx = File Systemx

What does *x* represent?

Your choice
for
your workload

What do we mean by *your choice*?

Launch, run, and scale . . .

feature-rich, high-performance, cost-effective, and fully managed storage . . .

Amazon FSx for NetApp ONTAP

Amazon FSx for OpenZFS

Amazon FSx for Windows File Server

Amazon FSx for Lustre

powered by **your choice** of the world's most popular file systems

Amazon FSx family

Amazon FSx for
Windows File Server

Amazon FSx for
NetApp ONTAP

Amazon FSx for
OpenZFS

Amazon FSx for
Lustre

Amazon FSx for Windows File Server

Amazon FSx for
Windows File Server

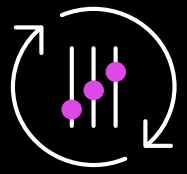

Fully managed

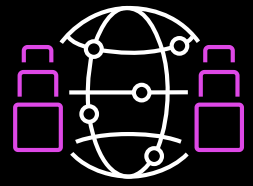

Single-AZ or
Multi-AZ

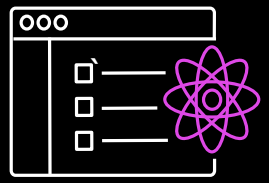

Built on
Windows Server

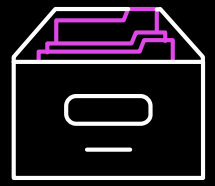

Snapshots and fully
managed backups

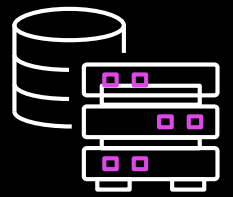

Multiple storage
type and
performance
options

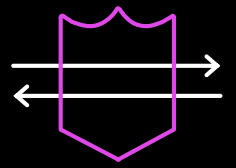

Integrates with
Microsoft Active
Directory

aws

FSx for Windows File Server use cases at a glance

Corporate IT applications

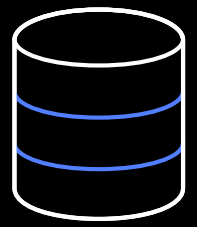

Microsoft SQL Server

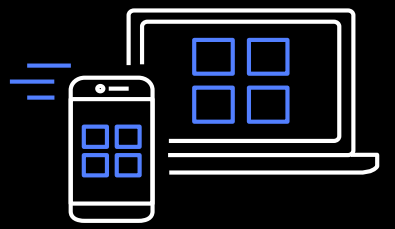

User shares

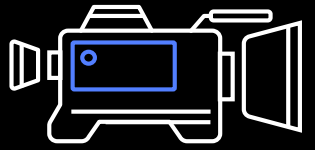

Media editing and transcoding

Customer reference

John Holland actively migrated more than **90 key applications** to FSx for Windows File Server **within 48 hours**

"**Amazon FSx for Windows File Server hasn't missed a beat**. The speed, the stability, and the ability to tune throughput speed and availability to our environment are quite positive for us."

Kier Morrison
GM of Information Technology Operations, John Holland Group

https://aws.amazon.com/solutions/case-studies/jhg-case-study/

Performance – Throughput capacity

Higher levels of throughput → higher levels of IOPS, and memory for caching

Throughput capacity Baseline throughput (MBps)	**Burst throughput** (MBps)
8	192
16	192
32	192
64	256
128	438
256	438
512	N/A
1,024	N/A
2,048	N/A

You get even higher throughput with in-memory caching on the file server (600 MBps – 3 GBps)

Cost-effective Windows storage in the cloud

	Single-AZ (per GB-month)	Multi-AZ (per GB-month)
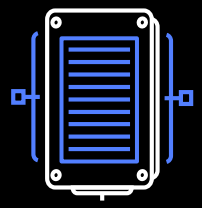 SSD storage TCO	6.5 cents	11.5 cents
 HDD storage TCO	0.65 cents	1.25 cents

Typical savings from deduplication and compression for general file shares is ~50%

Note: All pricing in USD, based on FSx for Windows File Server pricing in US East (N. Virginia)

High availability and data protection

Single-AZ

Continually monitors and addresses hardware failures

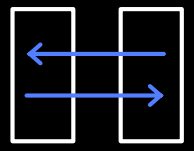

Replicates data within Availability Zone

Multi-AZ

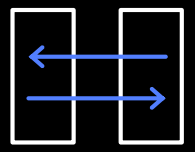

Replicates data *across* Availability Zones

Automatically fails over across Availability Zones

Snapshots

File-level restore (shadow copies)

Backups

Fully managed, automated backups to Amazon S3

aws

Amazon FSx for NetApp ONTAP

Amazon FSx for
Windows File Server

Amazon FSx for
NetApp ONTAP

Amazon FSx for
OpenZFS

Amazon FSx for
Lustre

Amazon FSx for NetApp ONTAP

THE FIRST, COMPLETE FULLY MANAGED ONTAP FILE SYSTEM IN THE CLOUD

Amazon FSx for NetApp ONTAP

Fully managed, PB+ scale

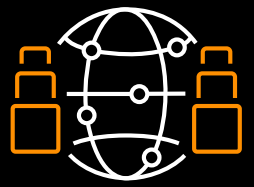

Single-AZ or Multi-AZ

High performance

Snapshot, clone, and replicate your data

Built-in storage efficiency and elastic tiering

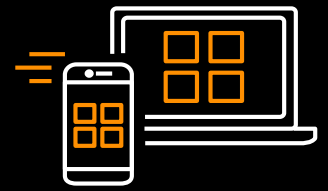

Accessible and multi-protocol (NFS, SMB, iSCSI)

Amazon FSx for NetApp ONTAP use cases at a glance

User and group file shares

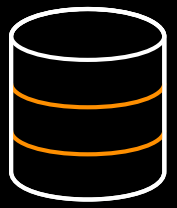

IT applications and databases

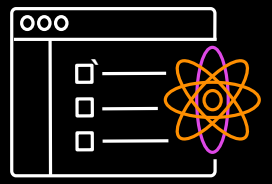

Line-of-business applications

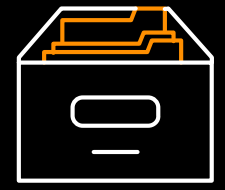

Backup, DR, and ransomware recovery

Amazon FSx for NetApp ONTAP: Performance, cost and scale

Storage tiers

SSD storage
- Sub-ms latencies
- Optimized for active data

Capacity pool storage
- Fully elastic (automatically grows/shrinks)
- Cost-optimized for infrequently accessed data

$0.03/GB-month*

Effective Multi-AZ storage cost

~20% of data

ONTAP's **built-in, automatic data tiering**
- Optimized for file workloads
- Intelligent and bidirectional

~80% of data

SSD performance for a fraction of the cost

*Pricing in USD, for US East (N. Virginia) Region. Assumes 80% of the files are infrequently accessed and 50% dedupe/compression

aws

Security and compliance

Data encrypted at rest and in transit

Integration with your organization's Active Directory

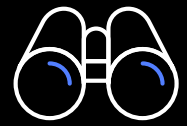

Monitoring and Logging

Monitor and log API calls with AWS CloudTrail and ONTAP audit logging

File access auditing

Supports auditing of access to files using ONTAP file access logging and FPolicy

Antivirus protection

Supports on-access anti-virus scanning using ONTAP vscan

Admin API access control and monitoring

Integrated with AWS IAM to control access to manage file systems

Network traffic access control

Integrated with Amazon VPC security groups to control network traffic

Industry-standard compliance eligibility

PCI-DSS, ISO, SOC, IRAP, GDPR, HIPAA

Customer reference

Pearson reduced RTO/RPO from hours to minutes

"As a fully managed service, FSx for NetApp ONTAP removed the operational overhead of our physical IT infrastructure freeing us to focus more on our application data and drive more value for our customers. With FSx for ONTAP, we **refreshed our data 10x faster**. We also **increased financial and HR job performance by 3x**."

Norm Nadeau
VP of Site Reliability Engineering, Pearson

Pearson Education

https://aws.amazon.com/solutions/case-studies/pearson-uses-amazon-fsx-for-netapp-ontap-video/

Amazon FSx for OpenZFS

Amazon FSx for
Windows File Server

Amazon FSx for
NetApp ONTAP

Amazon FSx for
OpenZFS

Amazon FSx for
Lustre

Amazon FSx for OpenZFS

Amazon FSx for OpenZFS

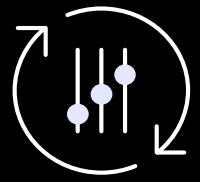

Fully managed

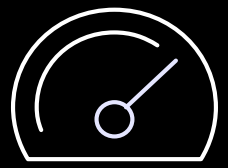

Optimized for high speed

Performance scaling with a few clicks

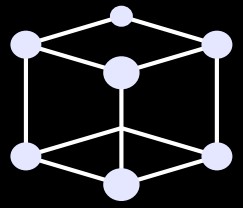

Advanced ZFS capabilities for working with data

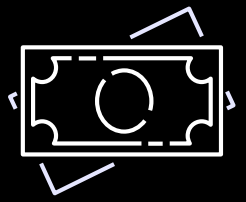

Low cost

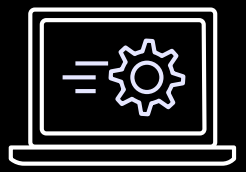

NFS v3, v4, v4.1, v4.2

FSx for OpenZFS use cases at a glance

High-performance network file system (NFS)-based workloads

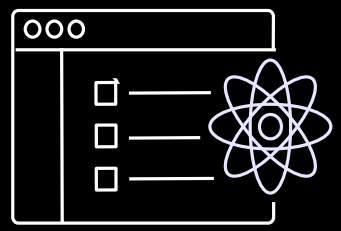

Engineering and frontend chip design

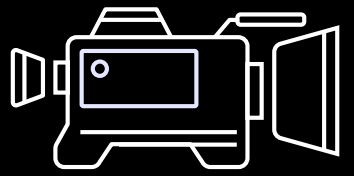

Media production

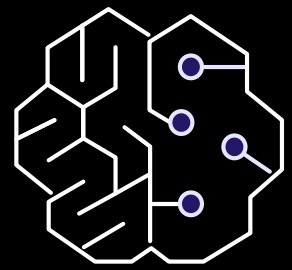

Genomics

Amazon FSx for OpenZFS performance

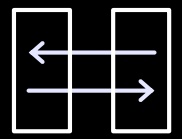

	Latency	IOPS	Throughput
What is it?	Average time to return the first byte of data	Number of operations (I/O) per second	Amount of data transferred per second
What can FSx for OpenZFS deliver?	Latencies of a **few hundred microseconds**	Up to **1 million IOPS**	Up to **12.5 GB/s**
When is it important?	**Sequential transaction-heavy applications like code repositories (Git) and content management systems (Wordpress)**		**Scalable data-intensive workloads like data and analytics, as well as ML**

Do more with less and adapt faster to changing business needs

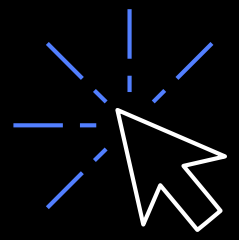

Storage and performance scaling in minutes*

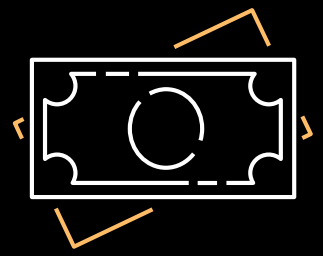

Zstandard compression (reduce storage usage by up to ~50%)

SSD storage
$0.09/GB-mo
($0.045/GB-mo w/ compression*)

Throughput
$0.26/MBps-mo
($0.13/MB-ps-mo w/ compression*)

*All Pricing in USD and assumes average compression savings of ~50% and is an effective price.

Protect your applications and data with robust security features

Automatic data encryption at rest and in transit

Automatic backups and maintenance

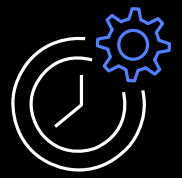

Automatic replication and recovery

Admin API access control & monitoring
(IAM, AWS CloudTrail)

Network traffic access control
(VPC security groups)

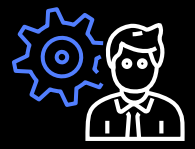

File-level permissions
(POSIX permissions and ACLs)

Industry-standard compliances
(PCI-DSS, ISO, SOC, IRAP, HIPAA)

Customer reference

Rev.com accelerated ML training for speech recognition and saved on storage costs

"With our self-managed NFS storage, our data scientists would encounter performance bottlenecks that slowed down their ability to develop and train models, and we often had to manually reconfigure our file servers to adapt to our changing needs. Amazon FSx for OpenZFS addresses these bottlenecks, and we can now scale our file system performance with the click of a button. We've eliminated the manual work of maintaining our own file servers, **while reducing costs by nearly 30%."**

Victor Yap
Senior MLOps Engineer, Rev.com

https://aws.amazon.com/fsx/openzfs/customers/#rev

Amazon FSx for Lustre

Amazon FSx for
Windows File Server

Amazon FSx for
NetApp ONTAP

Amazon FSx for
OpenZFS

Amazon FSx for
Lustre

Amazon FSx for Lustre

HIGH-PERFORMANCE FILE SYSTEM INTEGRATED WITH AMAZON S3

Amazon FSx for Lustre

Massively scalable performance

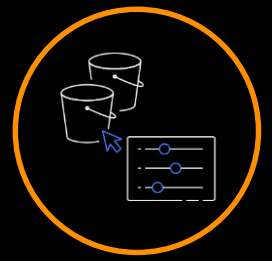

Seamless access to your data repositories

Fully managed

Native file system interface

Cost-optimized for compute-intensive workloads

Designed to be secure and compliant

Amazon FSx for Lustre use cases at a glance

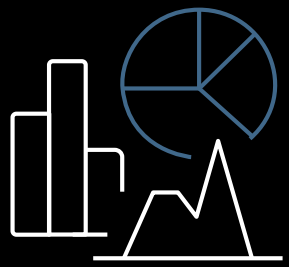

High performance computing

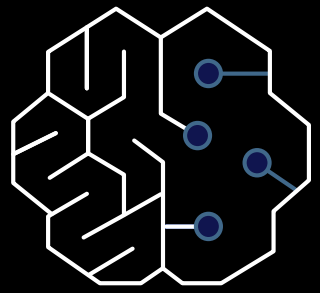

Machine learning

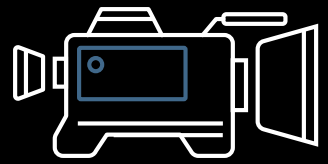

Media rendering and transcoding

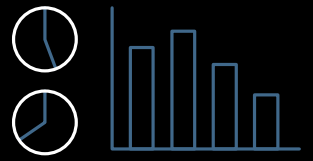

Analytics

FSx performance scales with storage capacity

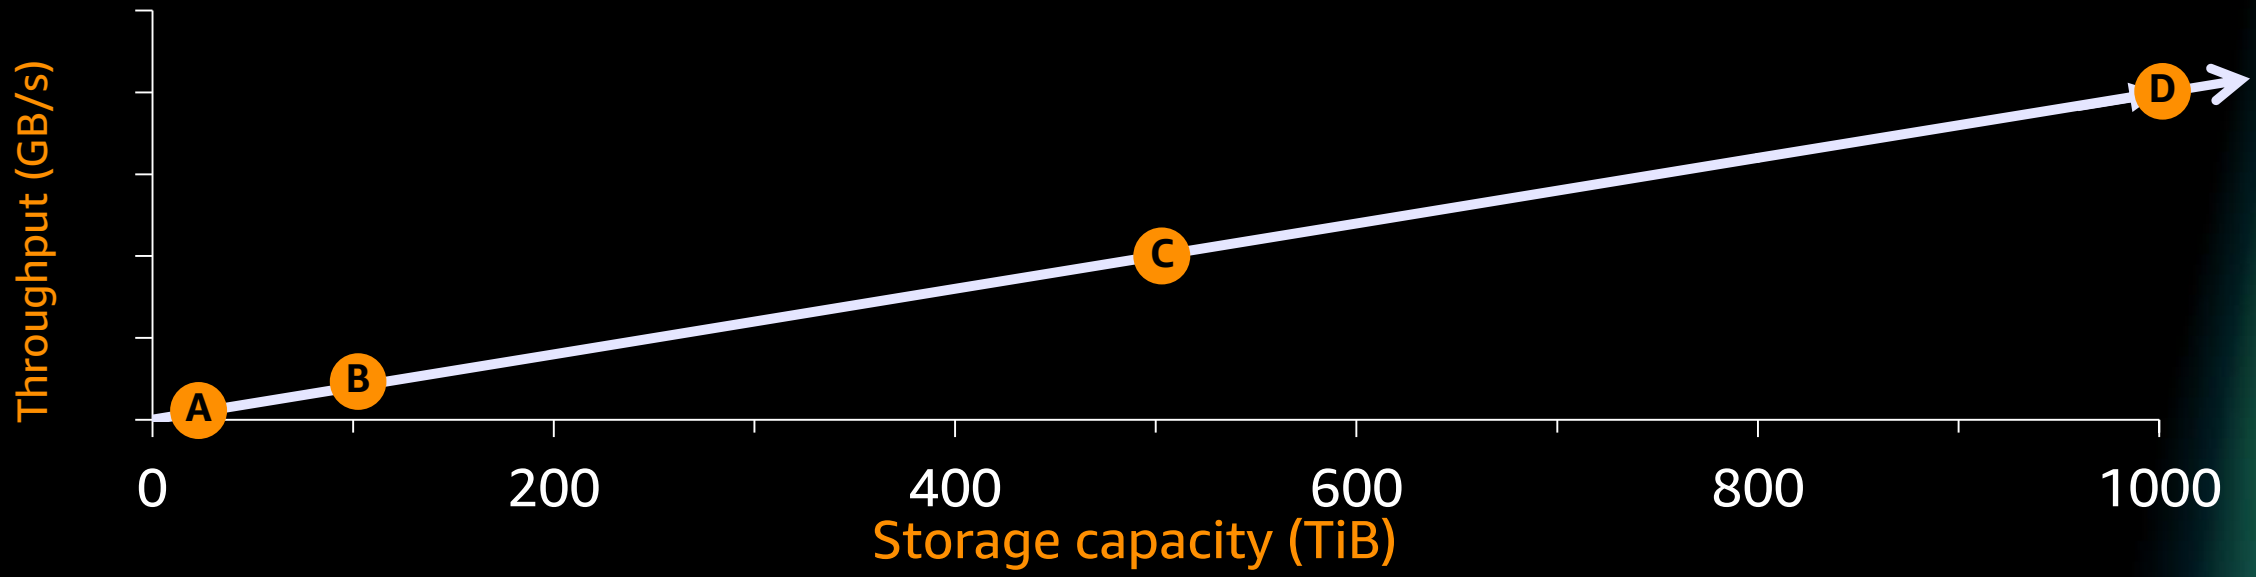

	Storage capacity	Baseline throughput	Burst throughput	IOPS	Latencies
A	10 TiB	2 GB/s	13 GB/s	Tens of thousands	Submillisecond
B	100 TiB	20 GB/s	130 GB/s	Hundreds of thousands	Submillisecond
C	500 TiB	100 GB/s	650 GB/s	>1 million	Submillisecond
D	1,000 TiB	200 GB/s	1,300 GB/s	Millions	Submillisecond

Lustre storage options optimized for price-performance

Storage type	Baseline throughput per TiB	Price per GB/s-hour
HDD Persistent	12 MB/s	$2.85
	40 MB/s	$2.85
NEW SSD Persistent	125 MB/s	$1.59
	250 MB/s	$1.15
	500 MB/s	$0.93
	1000 MB/s	$0.82

Amazon FSx for Lustre security deployment options

High and scalable performance

In all options, we support encryption at-rest and in-transit*

Customer reference

Using Amazon EC2 and Amazon FSx for Lustre allowed Joby to get faster results from their CFD workloads compared to on-premises high-performance computing infrastructure

Joby Aviation

"When we tried to run dozens of simulations at one time, we were reading and writing several gigabytes of data at a time, which slowed everything down. FSx for Lustre **eliminated** those **capacity problems**."

Alex Stoll
Aeromechanics Lead, Joby Aviation

https://aws.amazon.com/solutions/case-studies/joby-aviation-case-study/

Demo – Performance for stateful application on Amazon EKS with Amazon FSx for NetApp ONTAP

Demo scenarios – Performance for stateful application on Amazon EKS with Amazon FSx for NetApp ONTAP

- Stateful container application using Amazon EKS and Amazon FSx for NetApp ONTAP
- Provision Kubernetes persistent storage volumes using NetApp's Trident Container Storage Interface (CSI) driver
- FIO and IOping for running performance tests on FSxONTAP from within a K8S pod on Amazon EKS

Demo environment

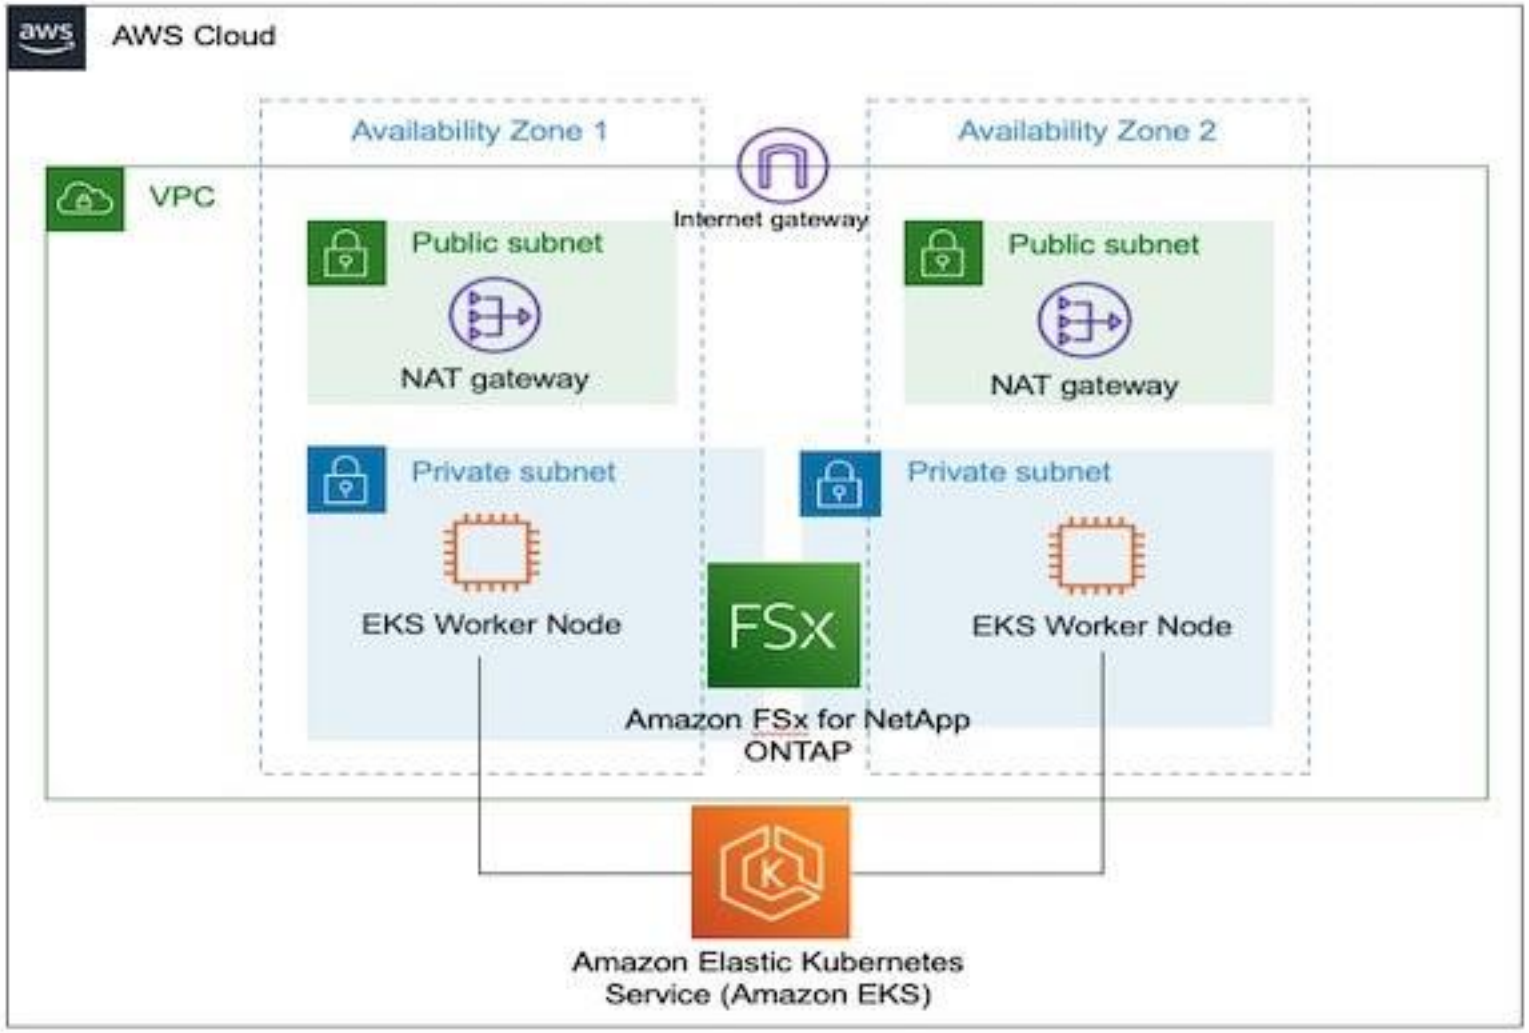

Performance summary

Scenario	Average IOPS (read)	Average IOPS (write)	Average throughput (read)	Average throughput (write)	Average latency
Same Availability Zone	32.7K	11.1K	137MB/s	45.2 MB/s	568.9 us
Different Availability Zone	26.8 K	9.1 K	112 MB/s	35.6 MB/s	1.26 ms

How do you *choose your* file system?

	Amazon FSx for NetApp ONTAP	Amazon FSx for Open ZFS	Amazon FSx for Windows File Server	Amazon FSx for Lustre
On-premises storage	NetApp, commodity NAS	ZFS or other Linux-based file servers	Windows file servers	Scale-out file storage (Lustre, GPFS, Isilon)
Unique Features	Multi-protocol, replication, cloning, intelligent tiering, file access auditing, snapshots, compression, deduplication	Lowest latency, scale-up NFS storage built on Graviton, compression, snapshots	Native Windows file sharing, file access auditing, snapshots, deduplication, AWS native on-premises caching	Scale-out performance, compression, S3 data processing capabilities
Use Cases	Enterprise IT, databases, line-of-business apps, test/dev, backup and DR (NetApp)	Enterprise IT, databases, line-of-business apps using NFS	Windows based user and group shares, Windows applications, SQL Server with HA	Machine Learning, HPC, media processing, data analytics, compute intensive applications

Next steps

Learn more about Amazon FSx:
https://aws.amazon.com/fsx/

Learn more about Amazon Storage Services:
https://aws.amazon.com/products/storage/

AWS Customer Success Stories:
https://aws.amazon.com/solutions/case-studies/

Amazon FSx workshop designed to help you better understand how to setup, configure, and get the best performance from an Amazon FSx file system:
https://github.com/aws-samples/amazon-fsx-workshop

Visit the AWS resource hub

Start building upon a scalable, reliable, and globally available infrastructure so that you can focus on innovation and bringing new applications to market. Dive deeper with these resources today.

- Accelerate innovation with AWS
- Get more performance for your applications at lower costs with AWS
- Global-scale solutions
- How startups succeed with AWS

https://tinyurl.com/for-every-app-hub-aws

Visit resource hub

AWS Training and Certification

Self – Paced Digital Training on AWS

Explore learning plans and 500+ digital courses from our new learning center, AWS Skill Builder, to help you achieve your goals on your schedule.

bit.ly/3lzVj0g

AWS Certification

Validate technical skills and cloud expertise to grow your career and business.

go.aws/3PwN3ff

aws

Thank you for attending AWS Innovate – For Every Application Edition

We hope you found it interesting! A kind reminder to **complete the survey.**
Let us know what you thought of today's event and how we can improve the event experience for you in the future.


aws-apj-marketing@amazon.com


twitter.com/AWSCloud


facebook.com/AmazonWebServices


youtube.com/user/AmazonWebServices


slideshare.net/AmazonWebServices


twitch.tv/aws

Thank you!