



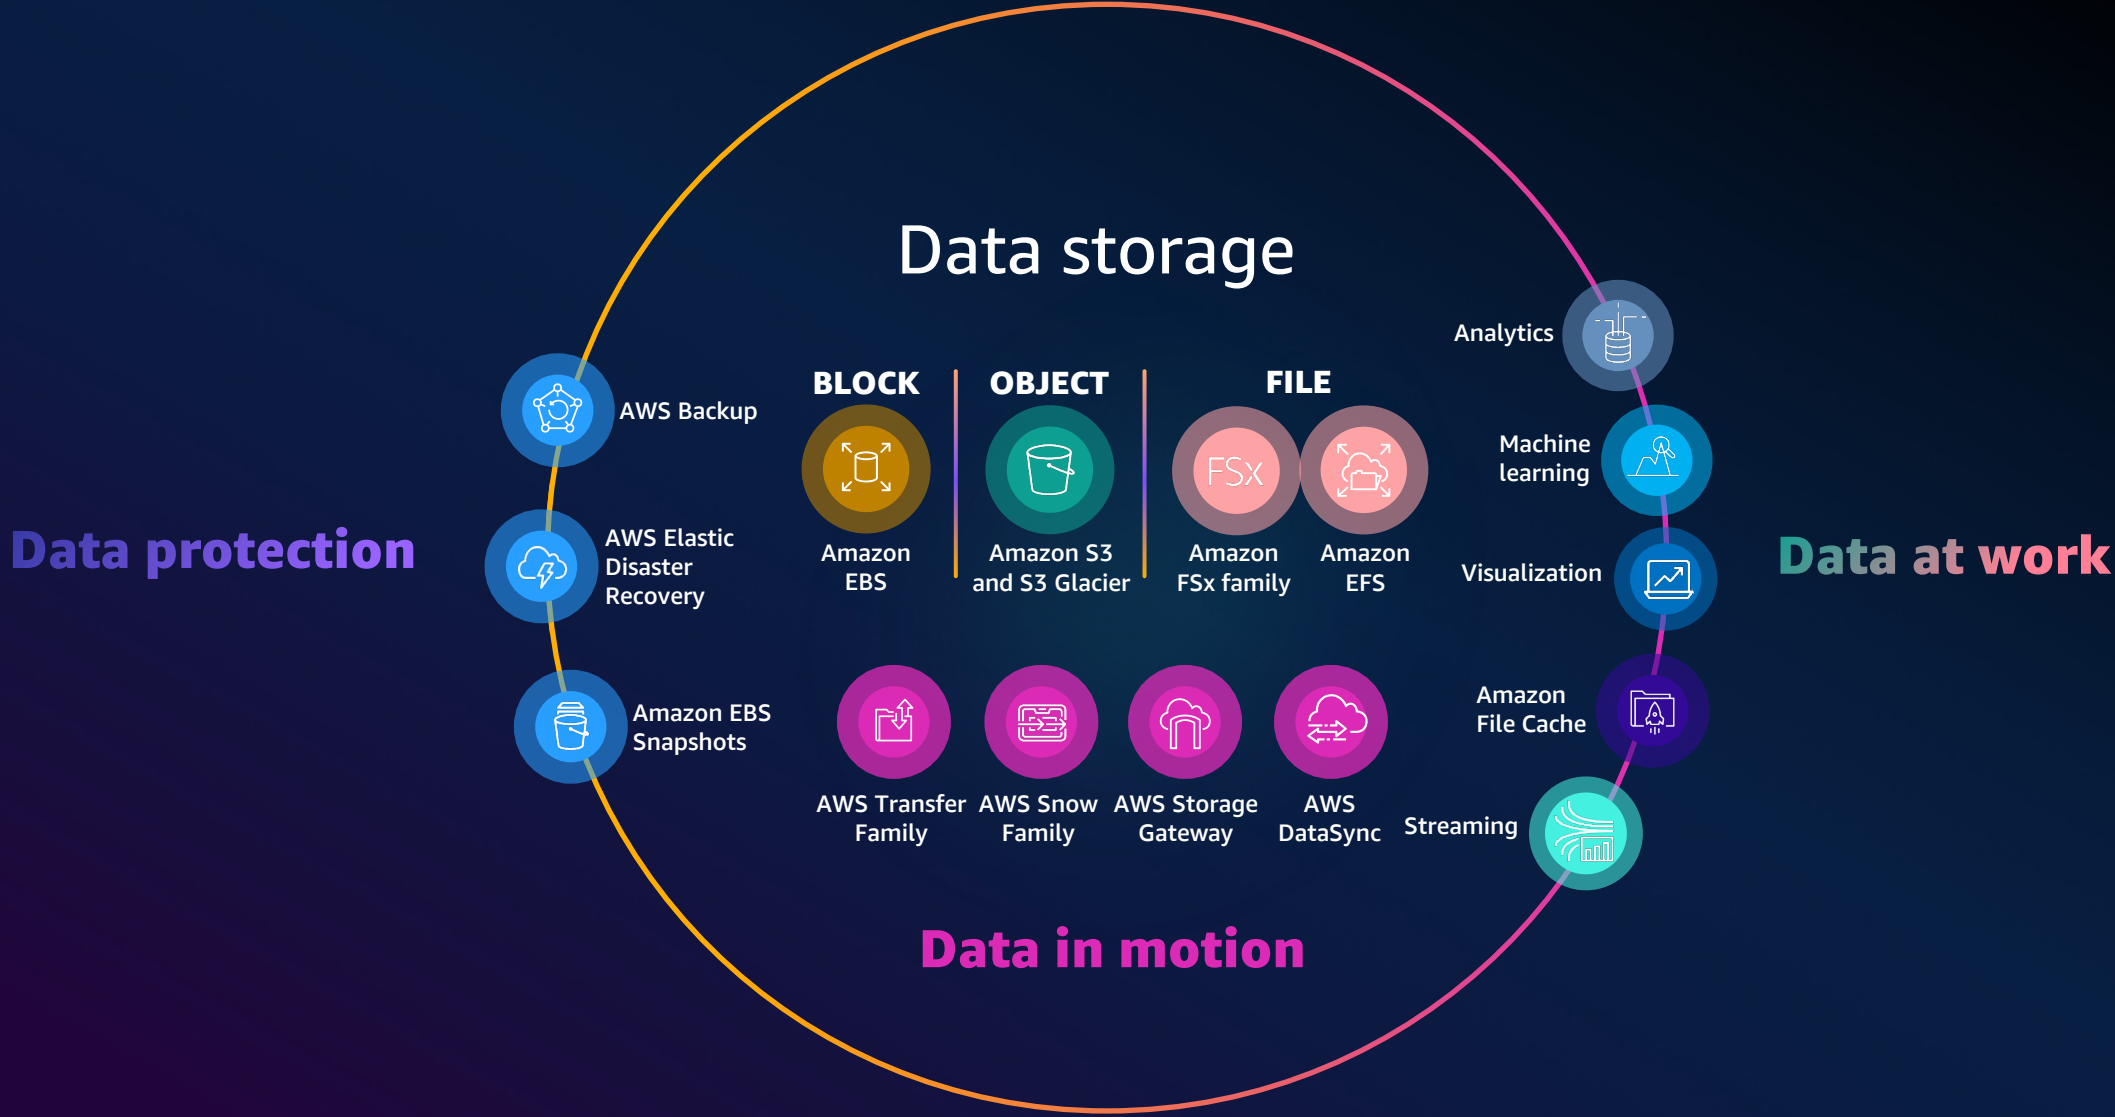
Managed services Lustre on AWS

Meriem Belhadj

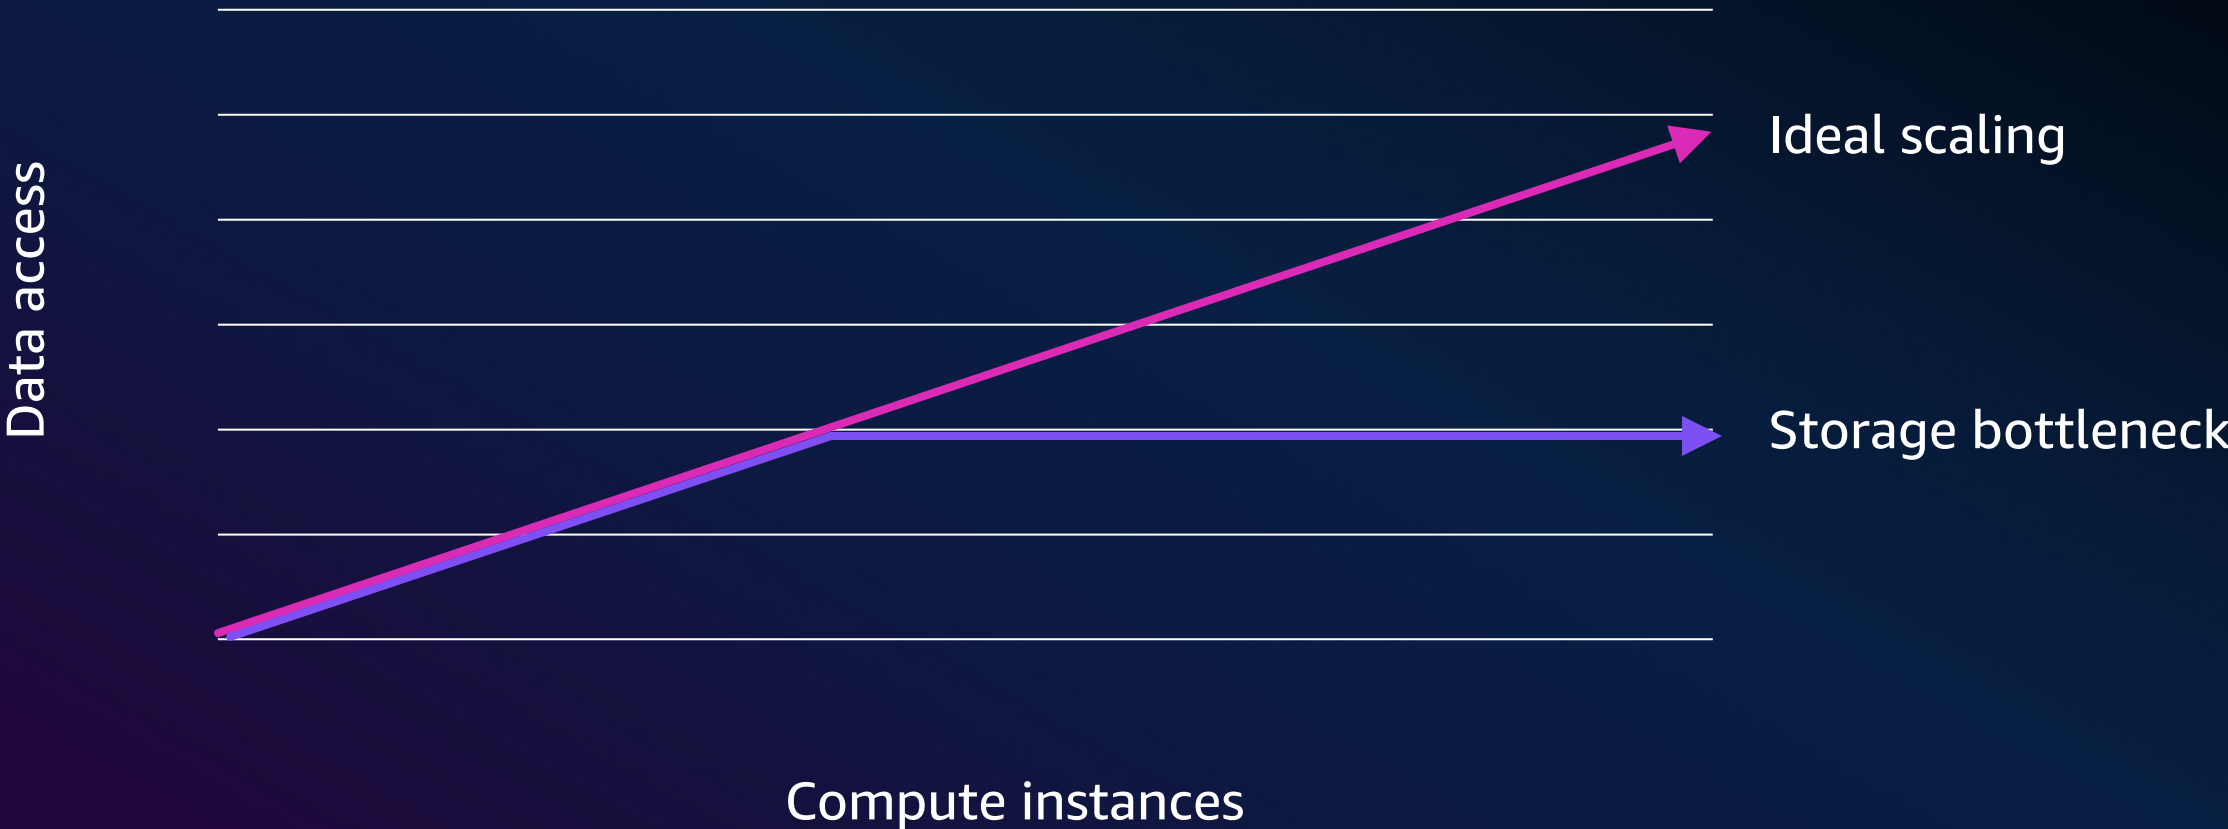
Sr. Storage Specialist Solutions Architect
Amazon Web Services



AWS cloud storage portfolio



Scalable compute requires scalable data access



FSx for Lustre: Designed for the most compute-intensive workloads

By industry



Financial services



Life sciences



Media and entertainment



Automotive



Semiconductor



Oil and gas

By application area



Big data analytics



Machine learning



High-performance computing

Amazon FSx for Lustre

The power of Lustre, fully-managed



Tested and operated
at an unprecedented
scale



Automatic infrastructure
monitoring and
replacement



Unified support
experience



Designed and tuned
to optimize AWS
resources

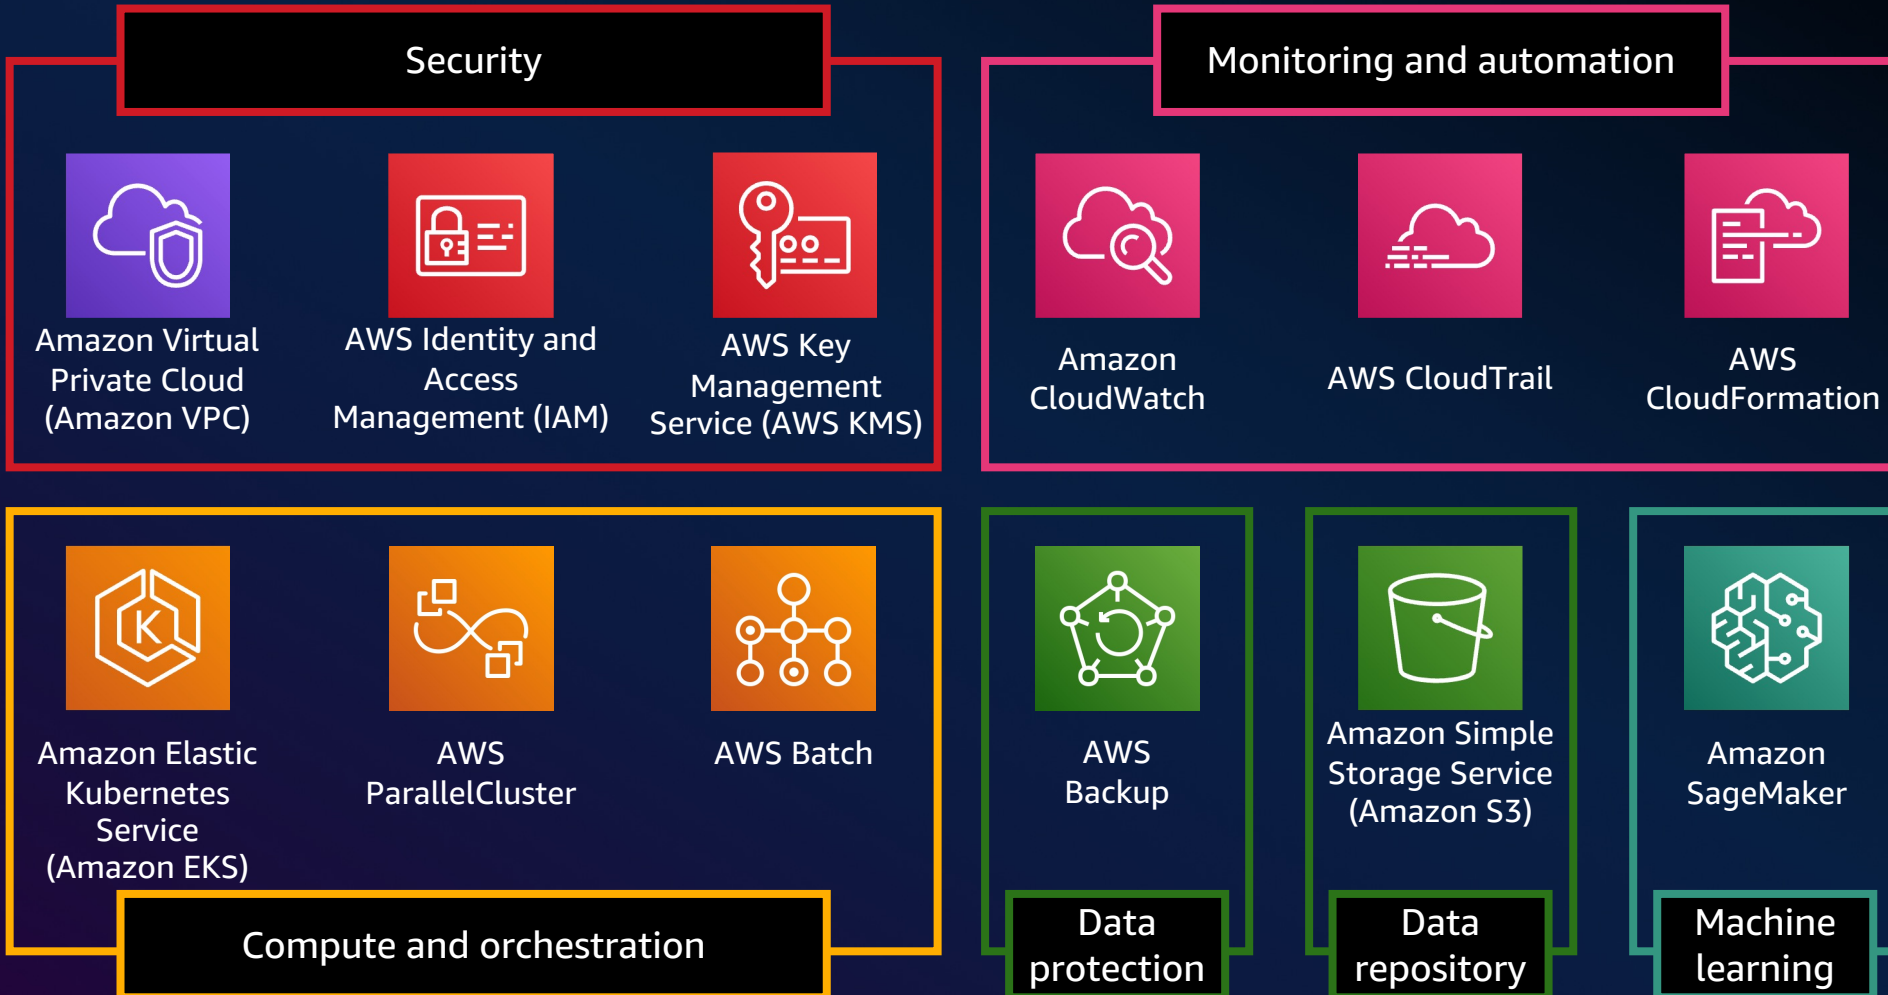


Automatic and
secure backups



API for infrastructure
as code


...and natively integrated into AWS



What FSx for Lustre brings to your workloads



Highly scalable
throughput
capacity



Storage options
optimized for
price performance



Integrate data lakes
through a fast file
interface

What FSx for Lustre brings to your workloads



Highly scalable
throughput
capacity

100s of GB/s
of throughput


Millions
of IOPS

Network bandwidth on Amazon EC2 instances is growing

Amazon Elastic Compute Cloud (Amazon EC2) instance bandwidth, Gbps

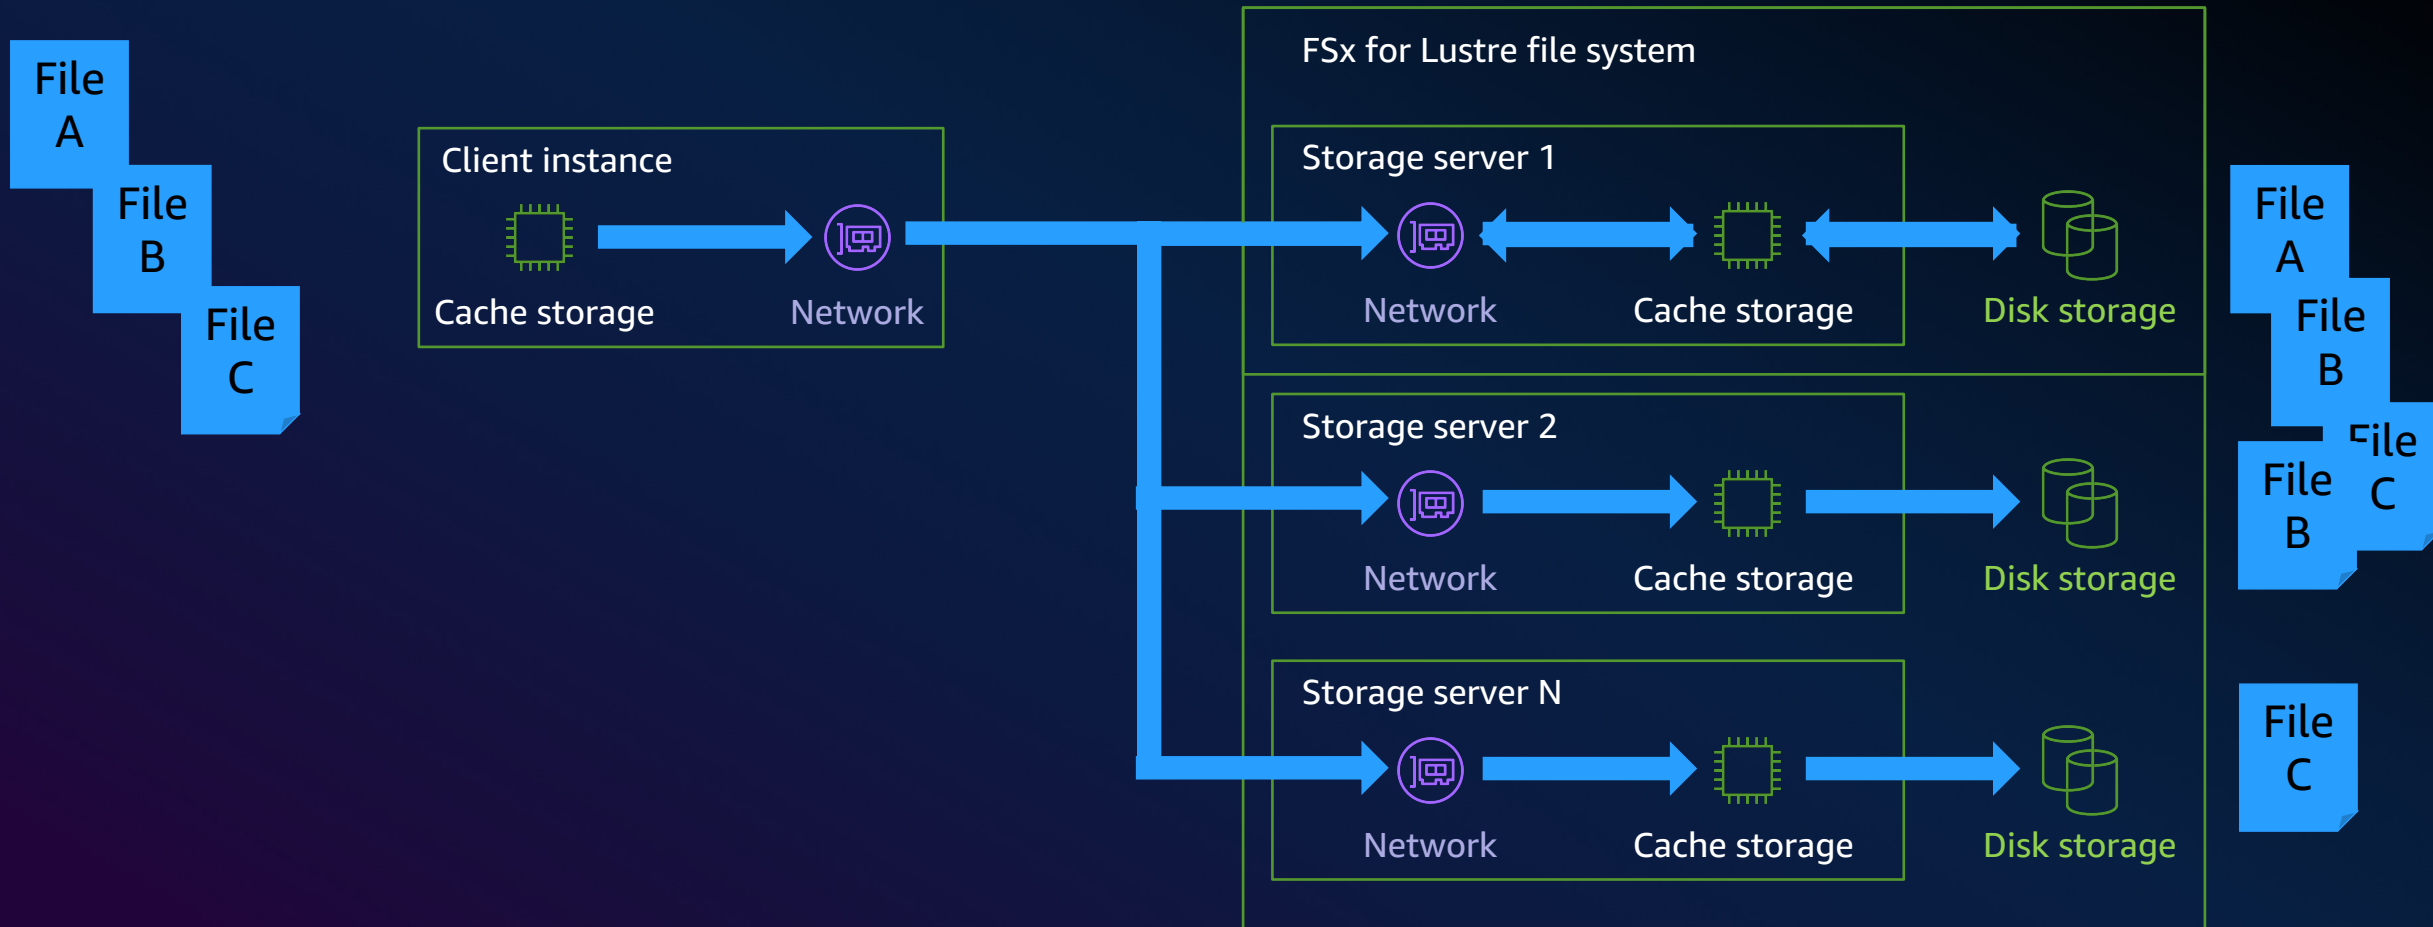


Lustre performance scales with storage capacity

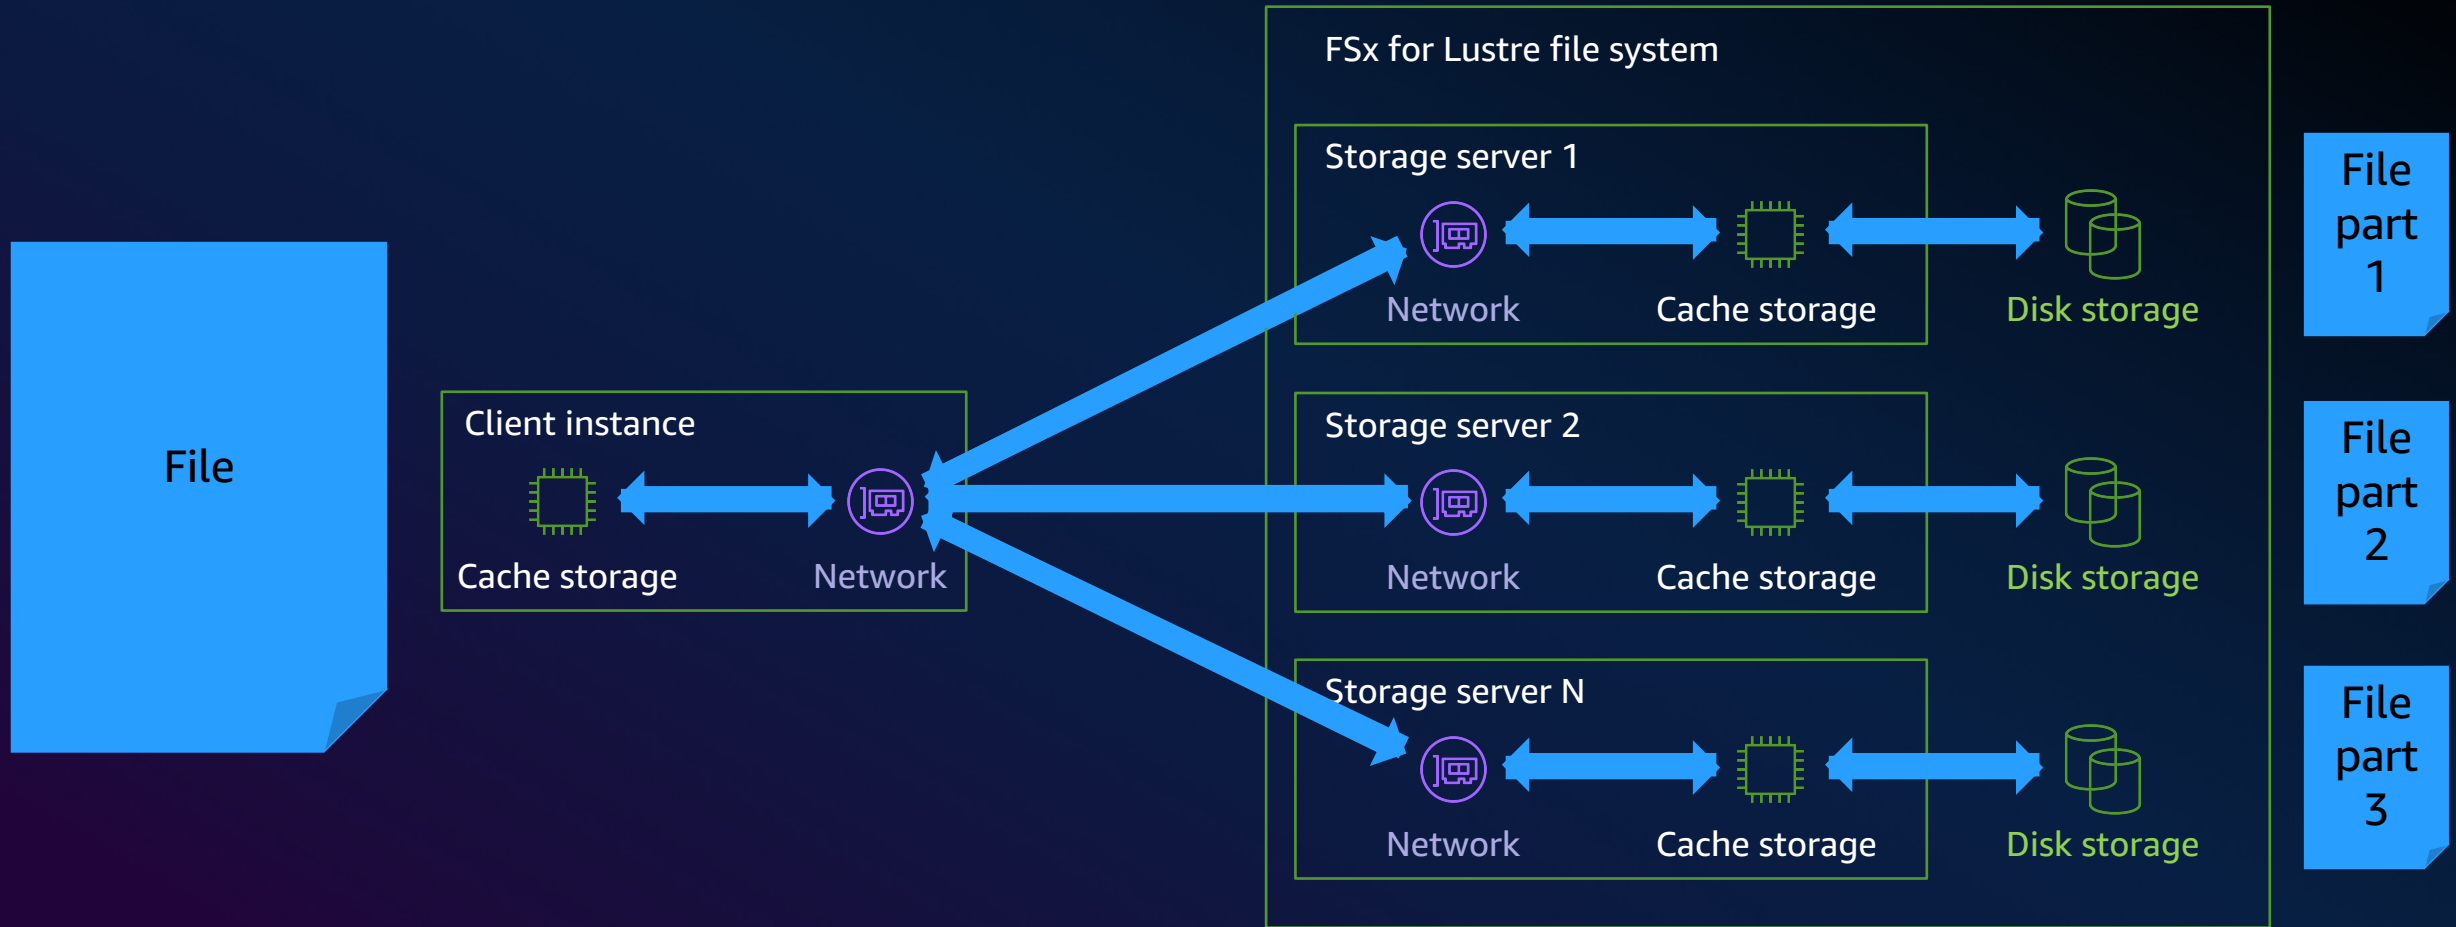


	Storage capacity	Baseline throughput	IOPS	Latencies
A	10 TiB	10 GB/s	Tens of thousands	Sub-millisecond
B	100 TiB	100 GB/s	Hundreds of thousands	
C	1,000 TiB	1,000 GB/s	Millions	

FSx for Lustre storage scaling and rebalancing



FSx for Lustre file IO



Lustre file IO

1 GB
file



>1 GB
file



1 TB
file





Shell augments their on-premises compute capacity by bursting to the cloud with Amazon EC2 clusters and Amazon FSx for Lustre.


This solution gives Shell the capability to **quickly scale up and down, and only purchase additional compute capacity when needed. With FSx for Lustre, Shell's GPU capacity is fully utilized.**



What FSx for Lustre brings to your workloads



Highly scalable
throughput
capacity



Storage options
optimized for
price performance



Integrate data lake
through a fast file
interface

Storage optimized for price performance



HDD-based persistent storage



SSD-based persistent storage



HDD-based persistent storage with an SSD cache


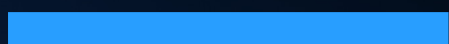












SSD-based scratch storage

Storage optimized for price performance

Storage type	Throughput per TiB	Price per unit of storage \$ per GB-month	
HDD persistent	12 MB/s	\$0.025	
	40 MB/s	\$0.083	
SSD persistent	125 MB/s	\$0.145	
	250 MB/s	\$0.210	
	500 MB/s	\$0.340	
	1,000 MB/s	\$0.600	

Storage optimized for price performance

Storage type	Throughput per TiB	Price per unit of storage \$ per GB-month	Price per unit of throughput, \$ per MB/s-month
HDD persistent	12 MB/s	\$0.025 	\$2.08 
	40 MB/s	\$0.083 	\$2.08 
SSD persistent	125 MB/s	\$0.145 	\$1.16 
	250 MB/s	\$0.210 	\$0.84 
	500 MB/s	\$0.340 	\$0.68 
	1,000 MB/s	\$0.600 	\$0.60 



GA 2023

Amazon FSx for Lustre


THROUGHPUT SCALING ON-DEMAND



INCREASED
AGILITY

- Scale your file system's throughput tier up or down without scaling storage capacity
- Easily accommodate the performance needs of growing compute clusters
- Manage periodic increases in data access in a cost-effective manner

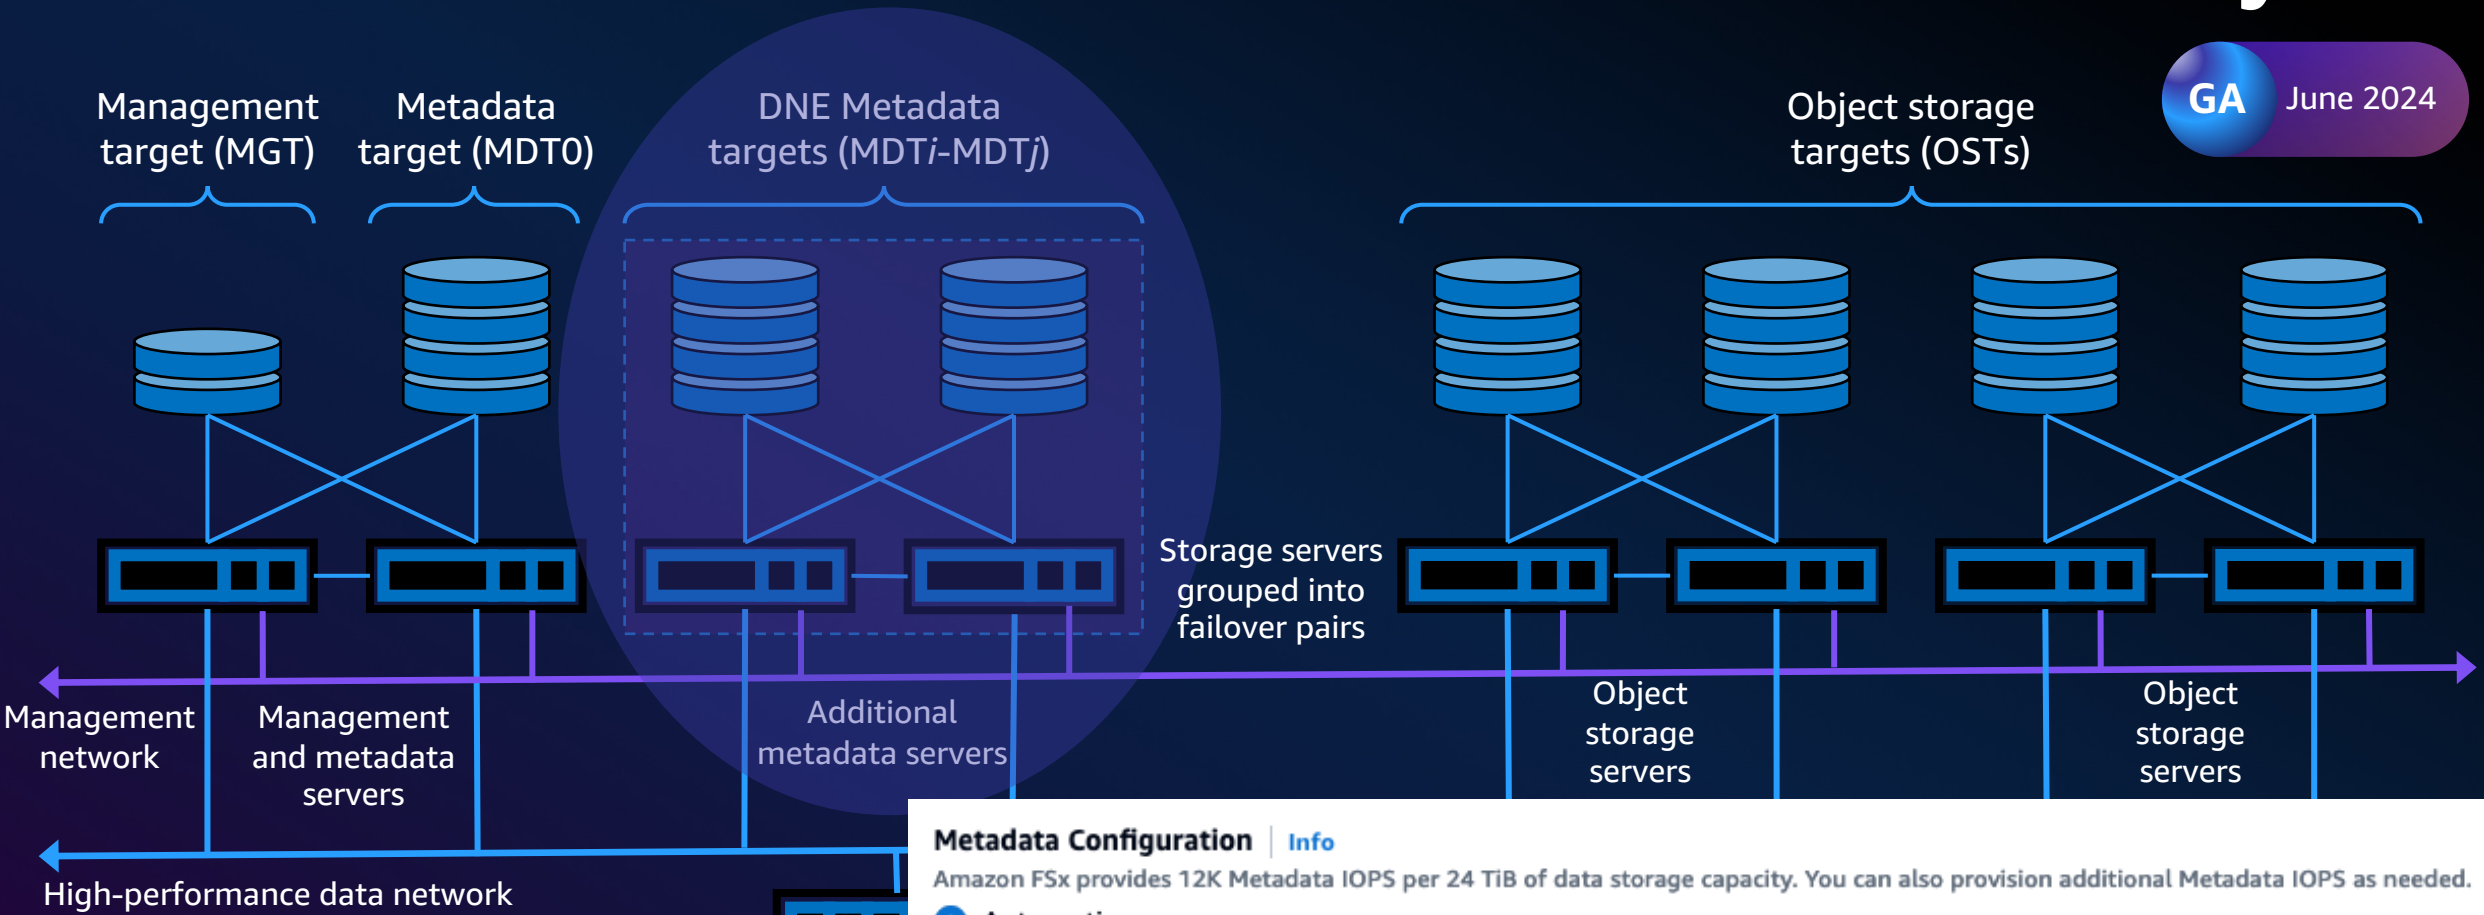
STORAGE SERVERS



STORAGE TARGETS
(DISKS)

New feature increases maximum metadata IOPS by 15x

GA June 2024



Metadata Configuration | [Info](#)

Amazon FSx provides 12K Metadata IOPS per 24 TiB of data storage capacity. You can also provision additional Metadata IOPS as needed.

- Automatic
12000 IOPS
- User-provisioned

Metadata IOPS

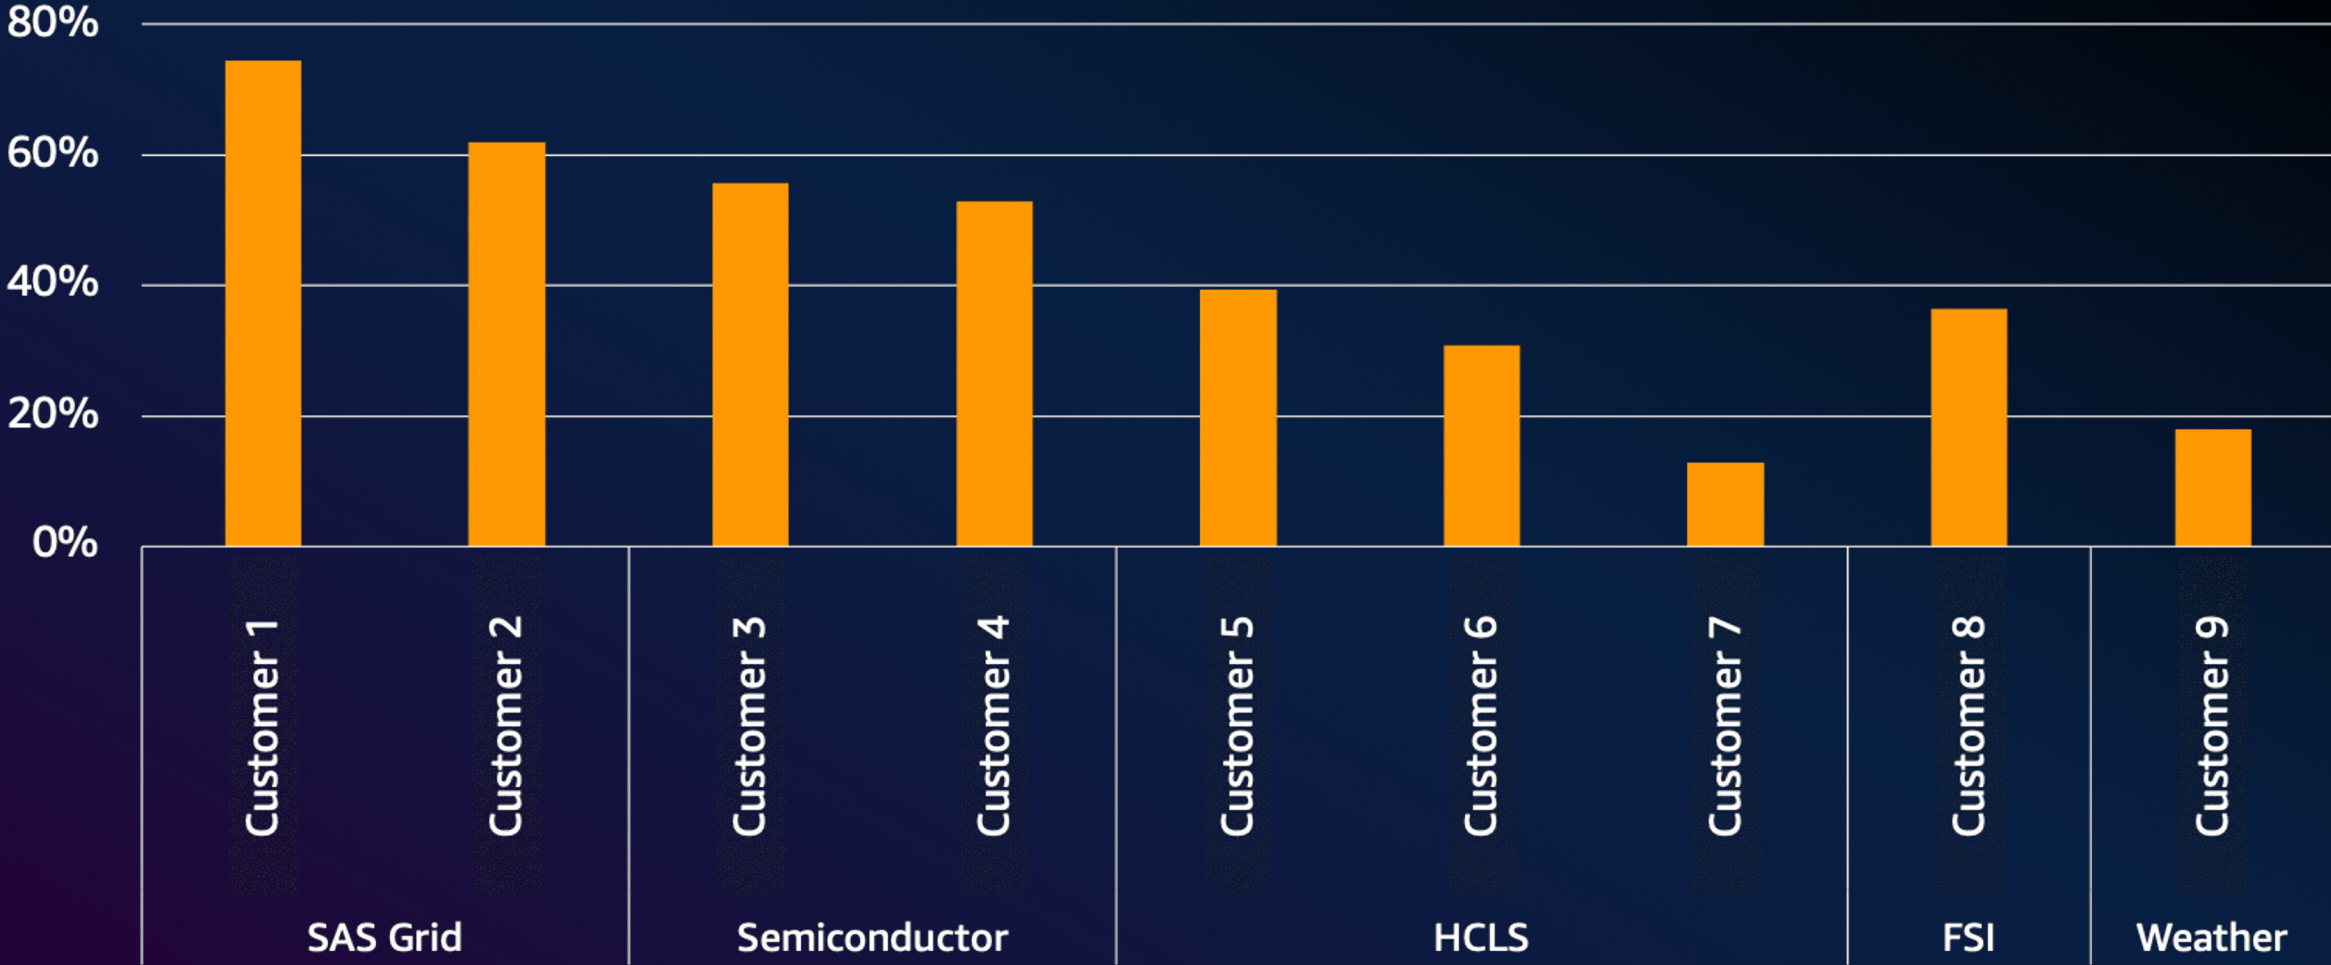
12000 IOPS

Supported sizes: 1,500, 3,000, 6,000 or increments of 12,000; Maximum 192,000 IOPS.

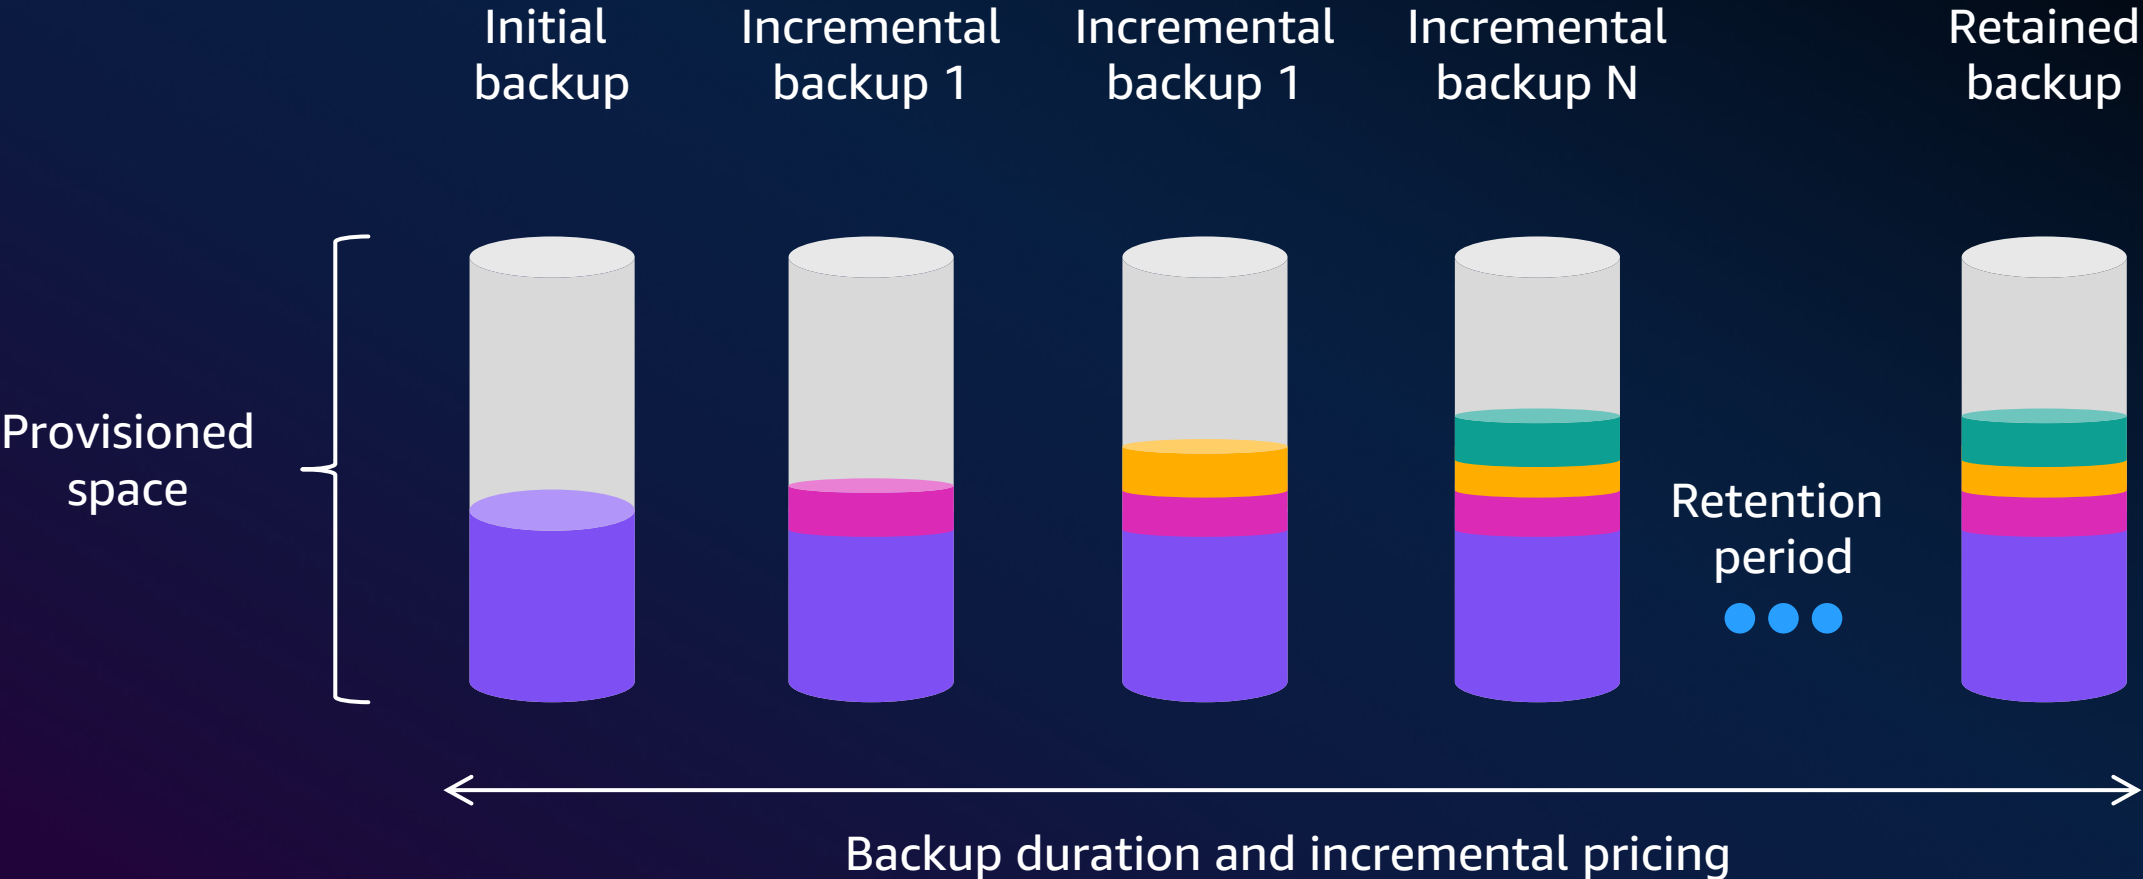
<https://docs.aws.amazon.com/fsx/latest/LustreGuide/managing-metadata-performance.html>



Data compression: Sample compression rates



Efficient backup storage



NETFLIX

Netflix accelerates large-scale ML training using Amazon FSx for Lustre to scale up and keep the data pipeline loaded

FSx for Lustre performance enables Netflix to saturate GPUs and to virtually eliminate GPU idle time

3–4x improvement using pre-compute and FSx for Lustre. ML model **training time reduced from a week to 1–2 days.**

FSx 

[Learn More
Amazon FSx for Lustre](#)



What FSx for Lustre brings to your workloads



Highly scalable
throughput
capacity



Storage options
optimized for
price performance



Integrate data lakes
through a fast file
interface

How FSx for Lustre enables HPC/ML with your data lake



Integrate your data lake with your ML/HPC compute through a fast file interface



AWS modern data architecture for machine learning (ML)

Pre-train

Fine-tune

Retrieval Augmented Generation (RAG)



Data sources

IOT/DEVICES

APP/LOGS

THIRD-PARTY DATA



Amazon S3



Amazon SageMaker



DIY/custom-built FMs



Amazon EC2



Amazon Bedrock



Amazon FSx for Lustre

Live or batch predictions (inference), model deployment, hosting & monitoring

PEOPLE


APPS

DEVICES



Reducing training time with FSx for Lustre

Amazon SageMaker: Patient classification using liver dataset





Toyota Research Institute chooses FSx for Lustre to reduce object recognition ML training times

“We needed a parallel file system for our ML training data sets and chose Amazon FSx for Lustre ... The **integration with AWS services, including S3**, also made it the preferred option for our high performance file storage.”

David Fluck, Software Engineer
Toyota Research Institute

FSx

[Learn More
Amazon FSx for Lustre](#)



Thank you!

Meriem Belhadj
mebelhad@amazon.fr



Cloud Storage
on AWS

