



aws INNOVATE

AI/ML EDITION

24 February 2022

Accelerate your SageMaker model training with Amazon FSx for Lustre and Amazon S3

Gaurav Singh

Senior Solutions Architect, ISV
AISPL

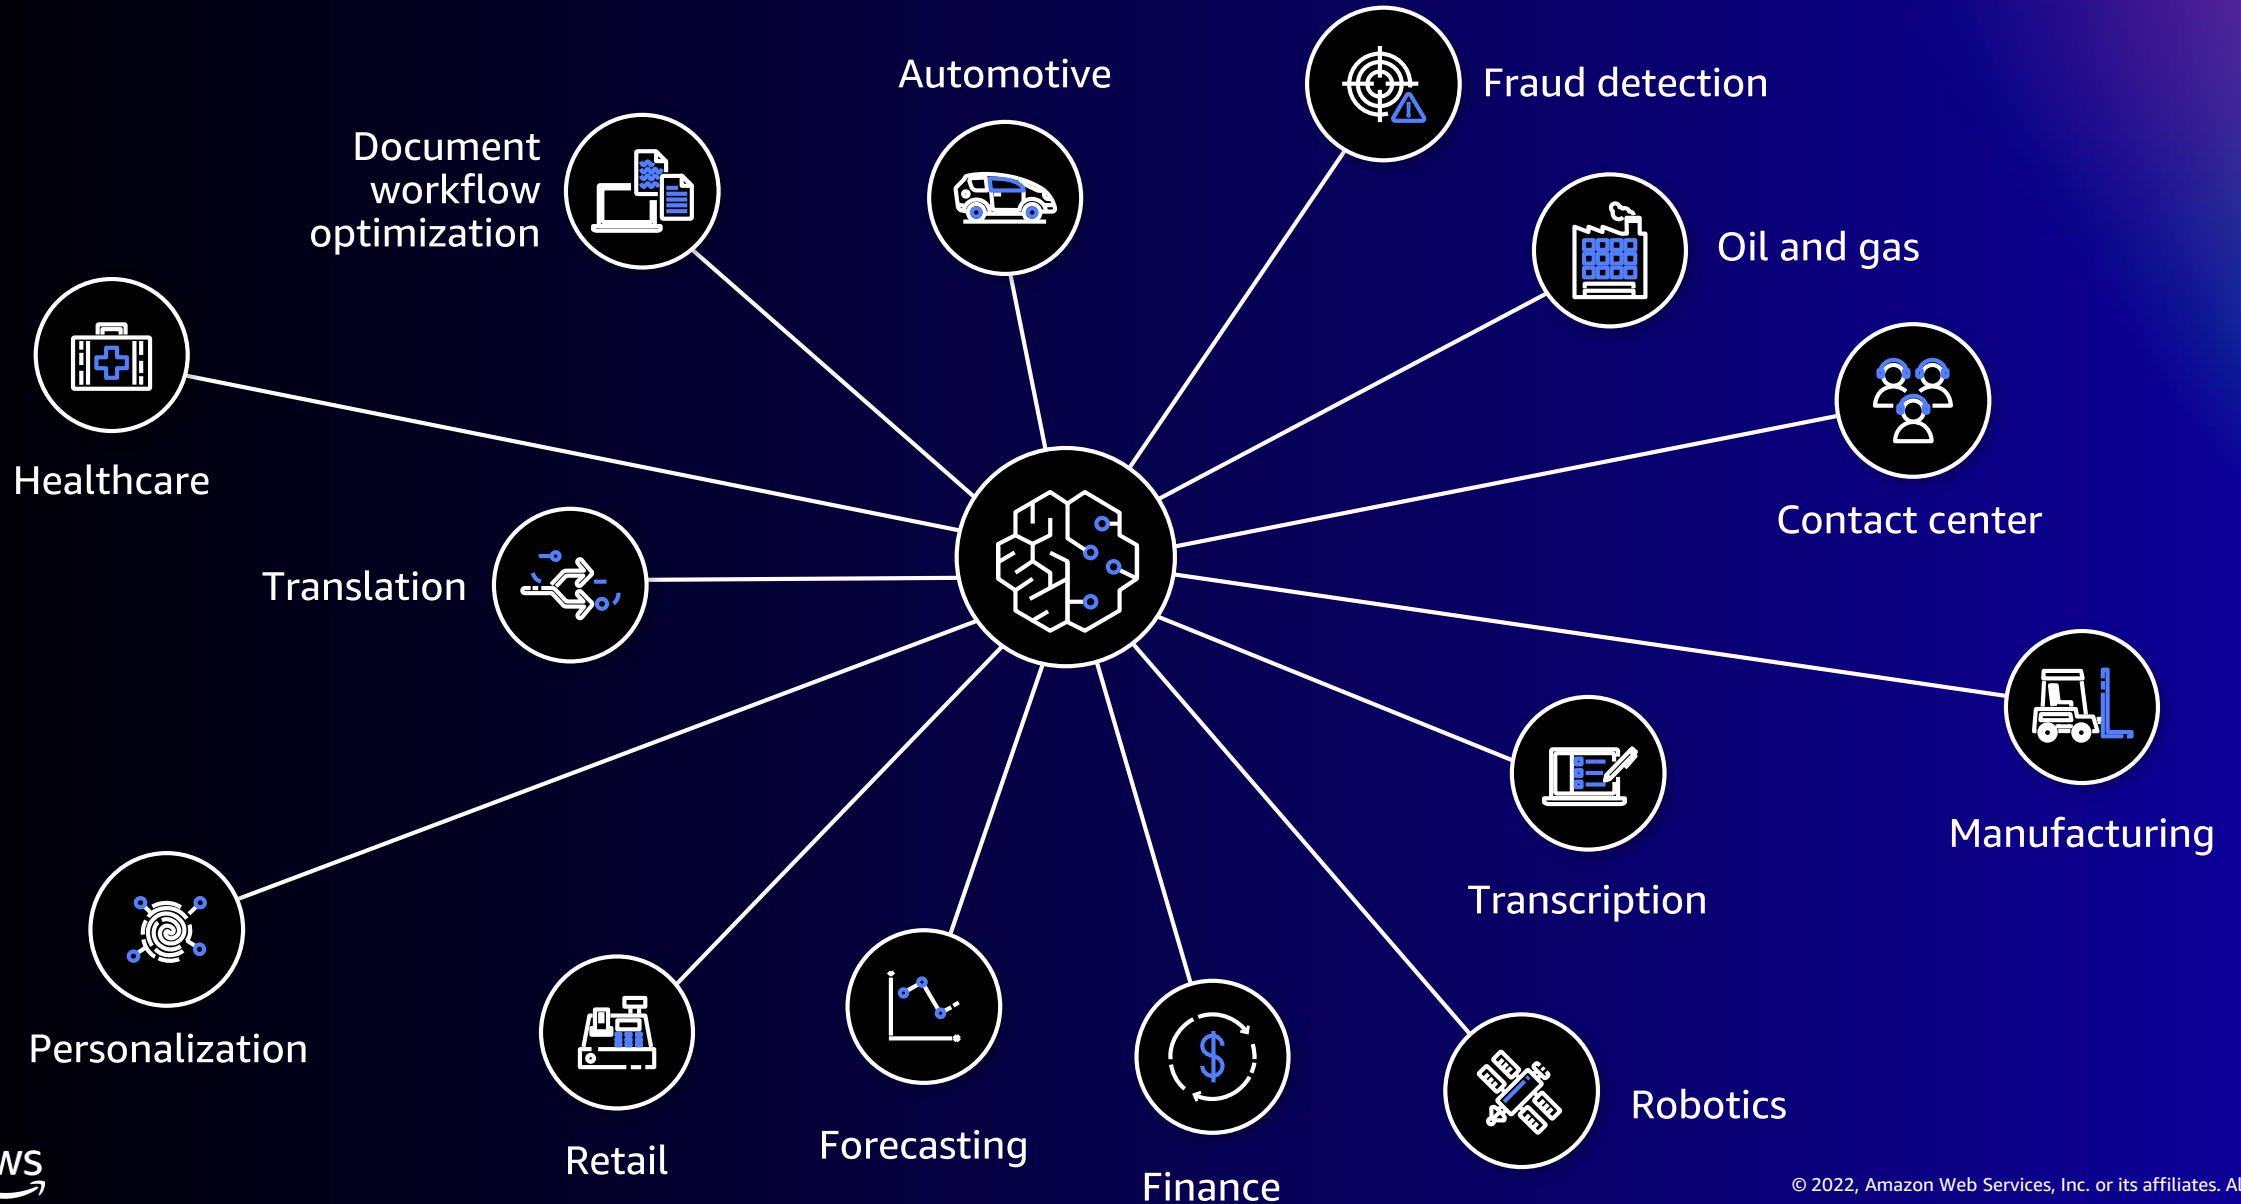


Agenda

1. AI/ML with AWS
2. Distributed training on Amazon SageMaker
3. Storage options with Amazon SageMaker Training
4. Demo
5. Observations from demo
6. Recap

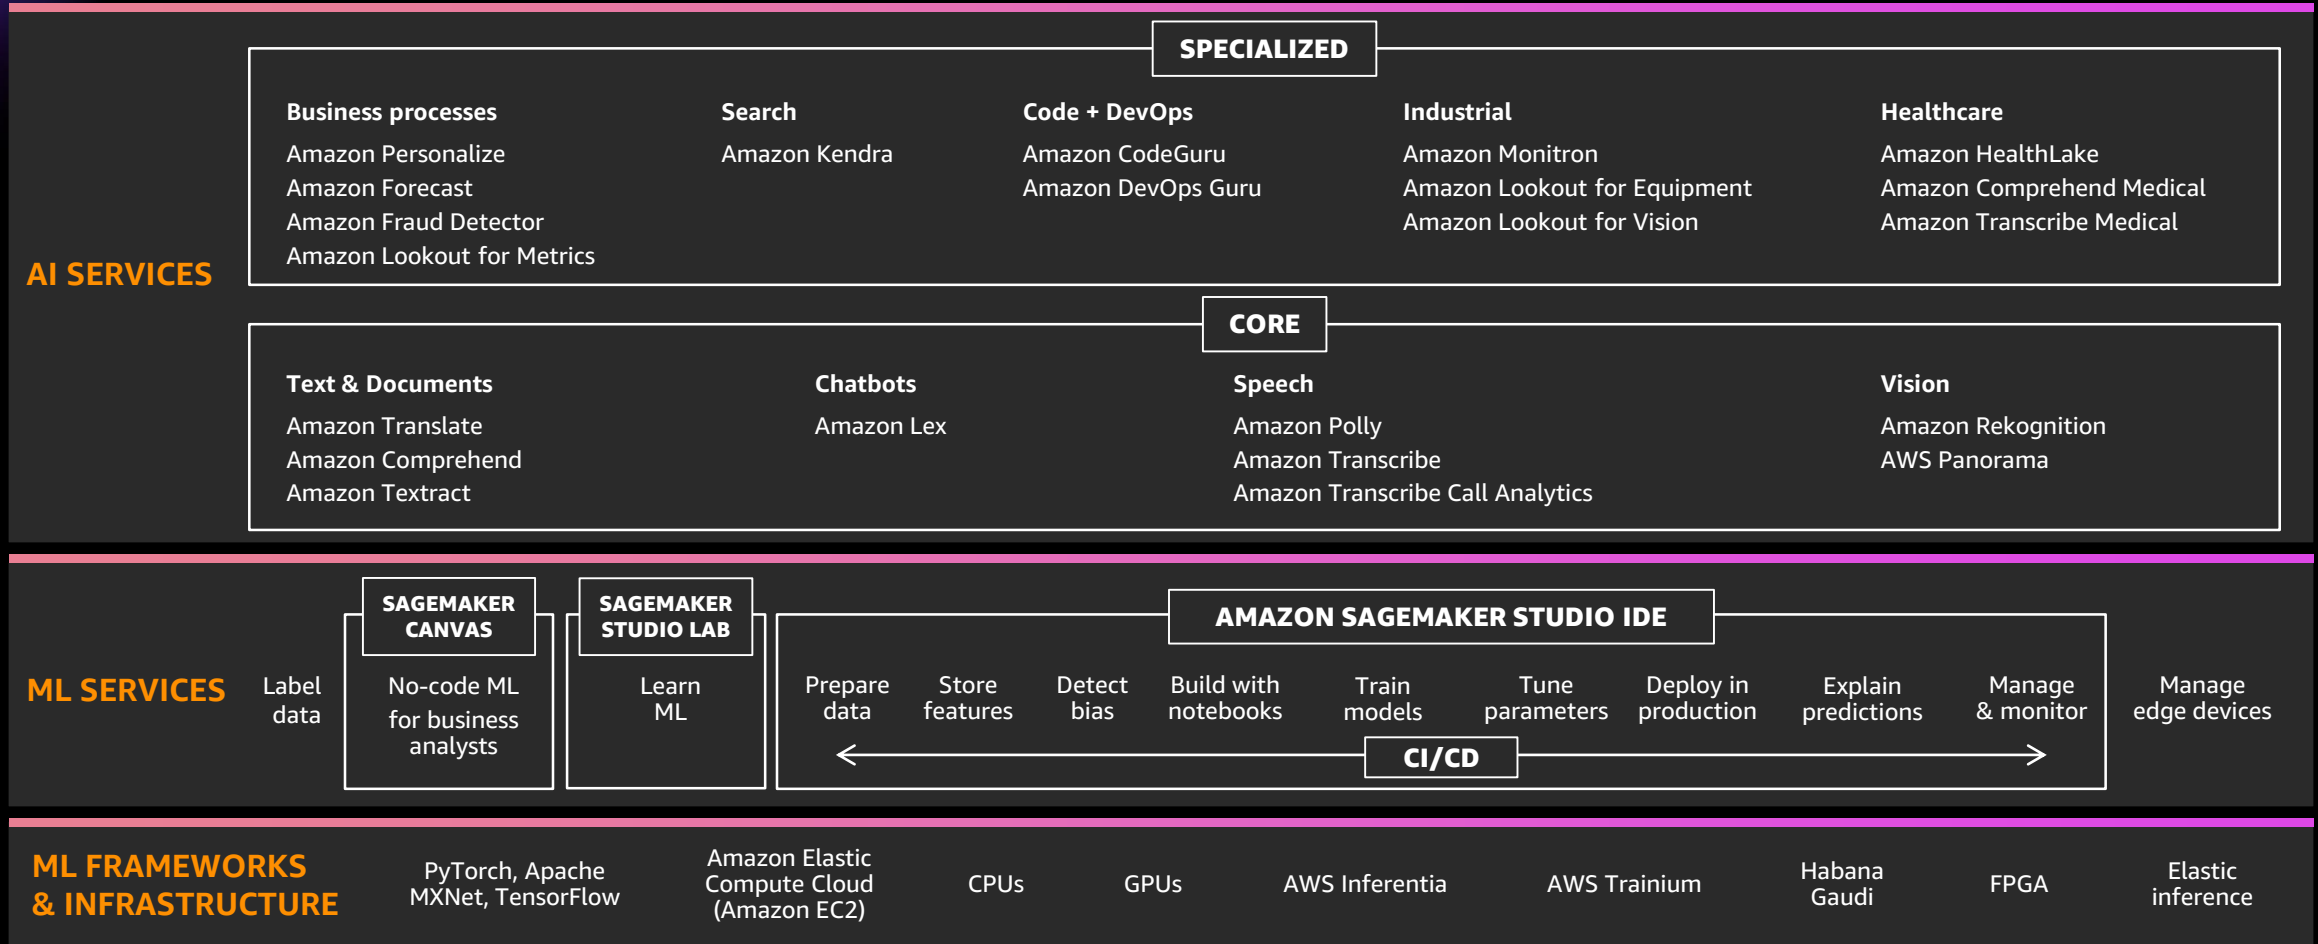
AI/ML with AWS

The reach of ML is growing

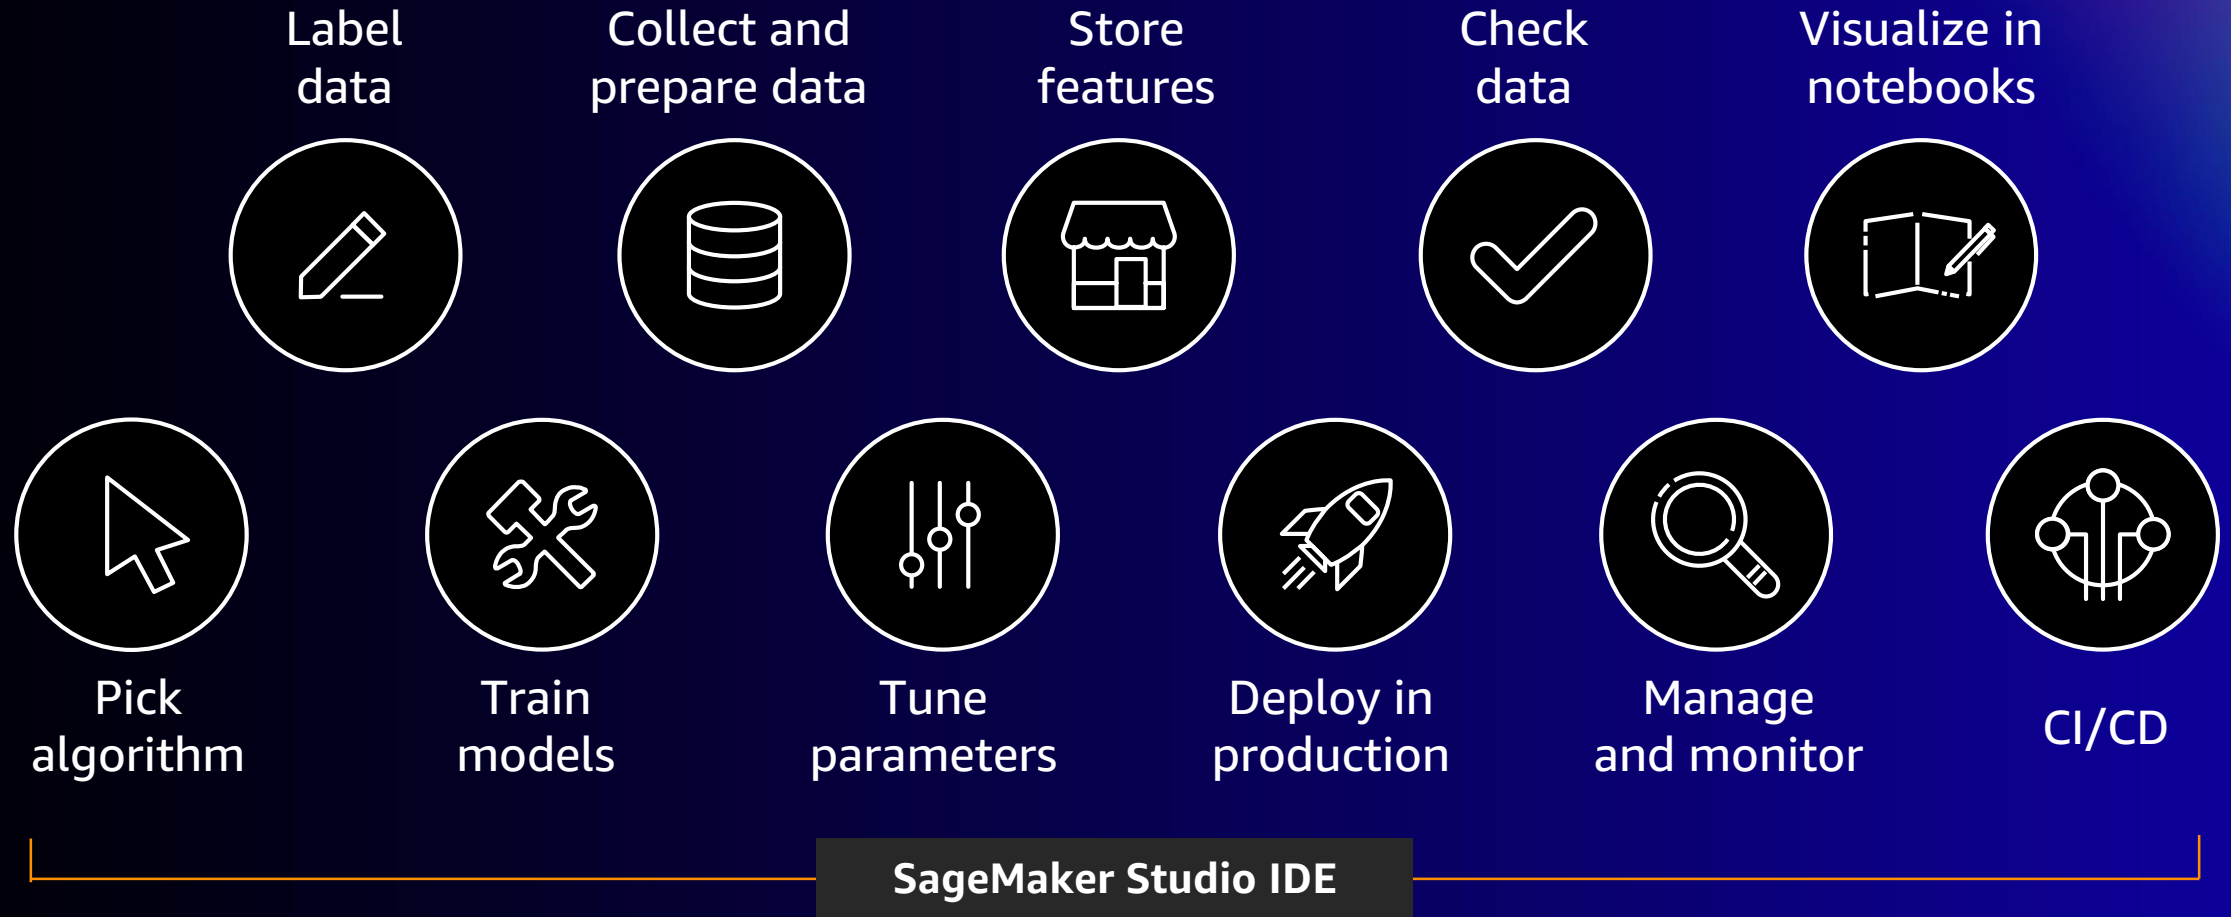


The AWS ML stack

BROADEST AND MOST COMPLETE SET OF MACHINE LEARNING CAPABILITIES

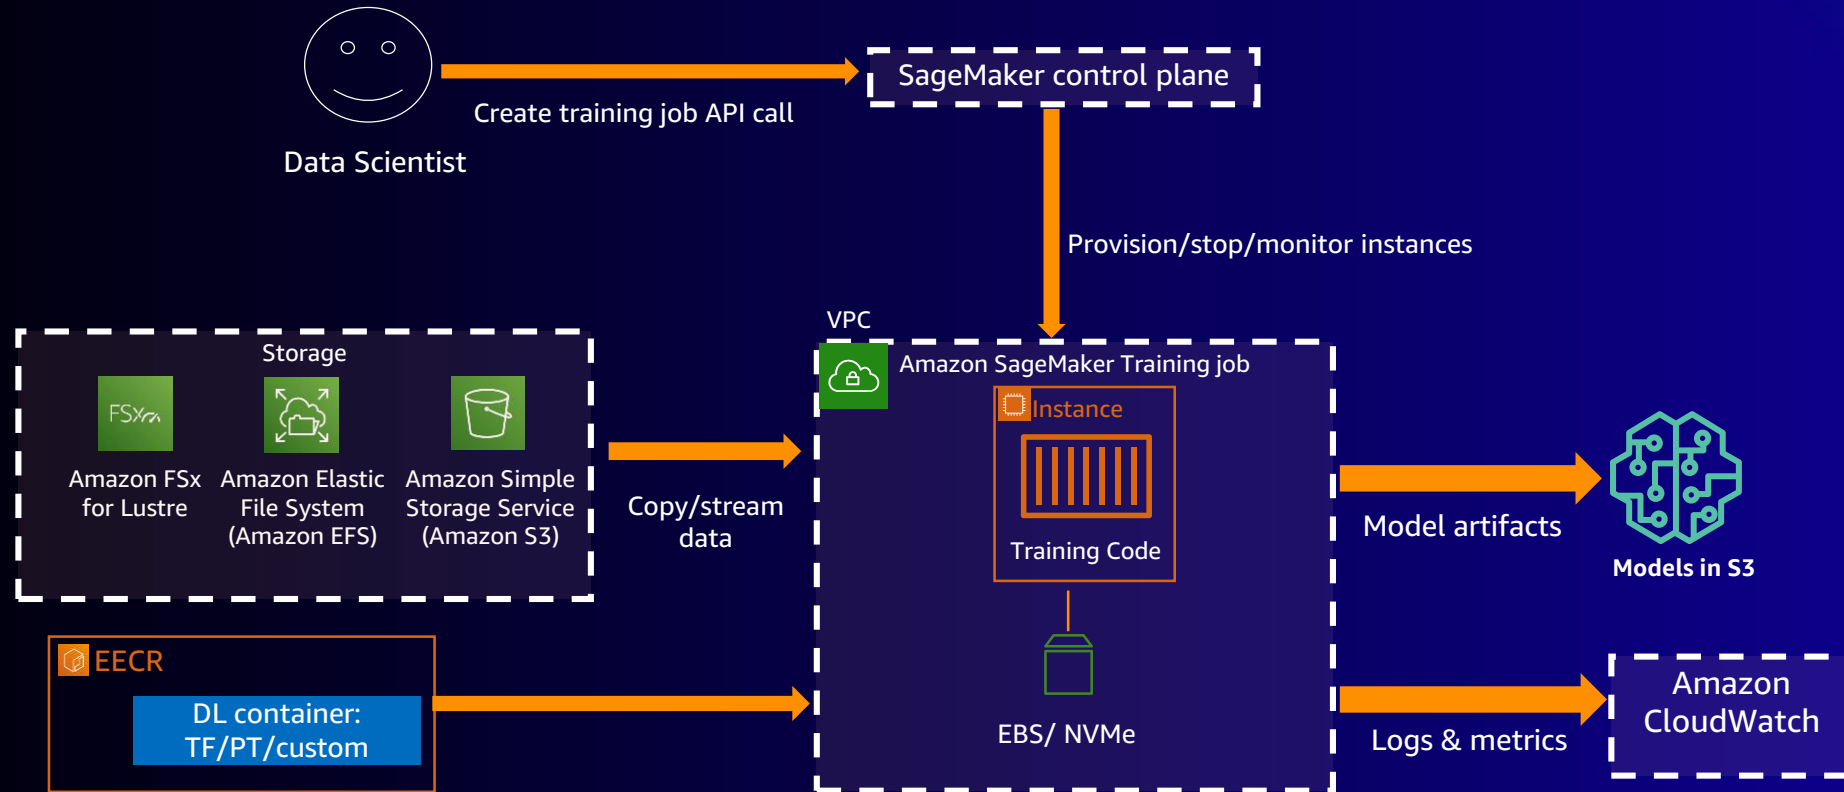


Amazon SageMaker: Built to make ML more accessible

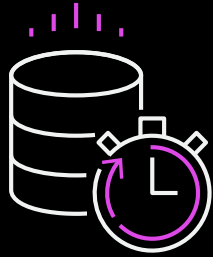


Distributed training on Amazon SageMaker

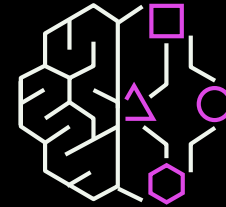
Amazon SageMaker Training – a refresher



Challenges with large datasets and models



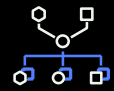
Large datasets take a long time to train, creating a bottleneck



Large models or models with many parameters can exceed the memory capacity of a single GPU

Model parallelism

Automated and efficient
model partitioning



Data parallelism

Reduced
training time



Minimal code
change



Distributed training on Amazon SageMaker

<https://aws.amazon.com/sagemaker/distributed-training/>

Optimized
for AWS network



Optimized for training
large models



Support for popular
ML framework APIs



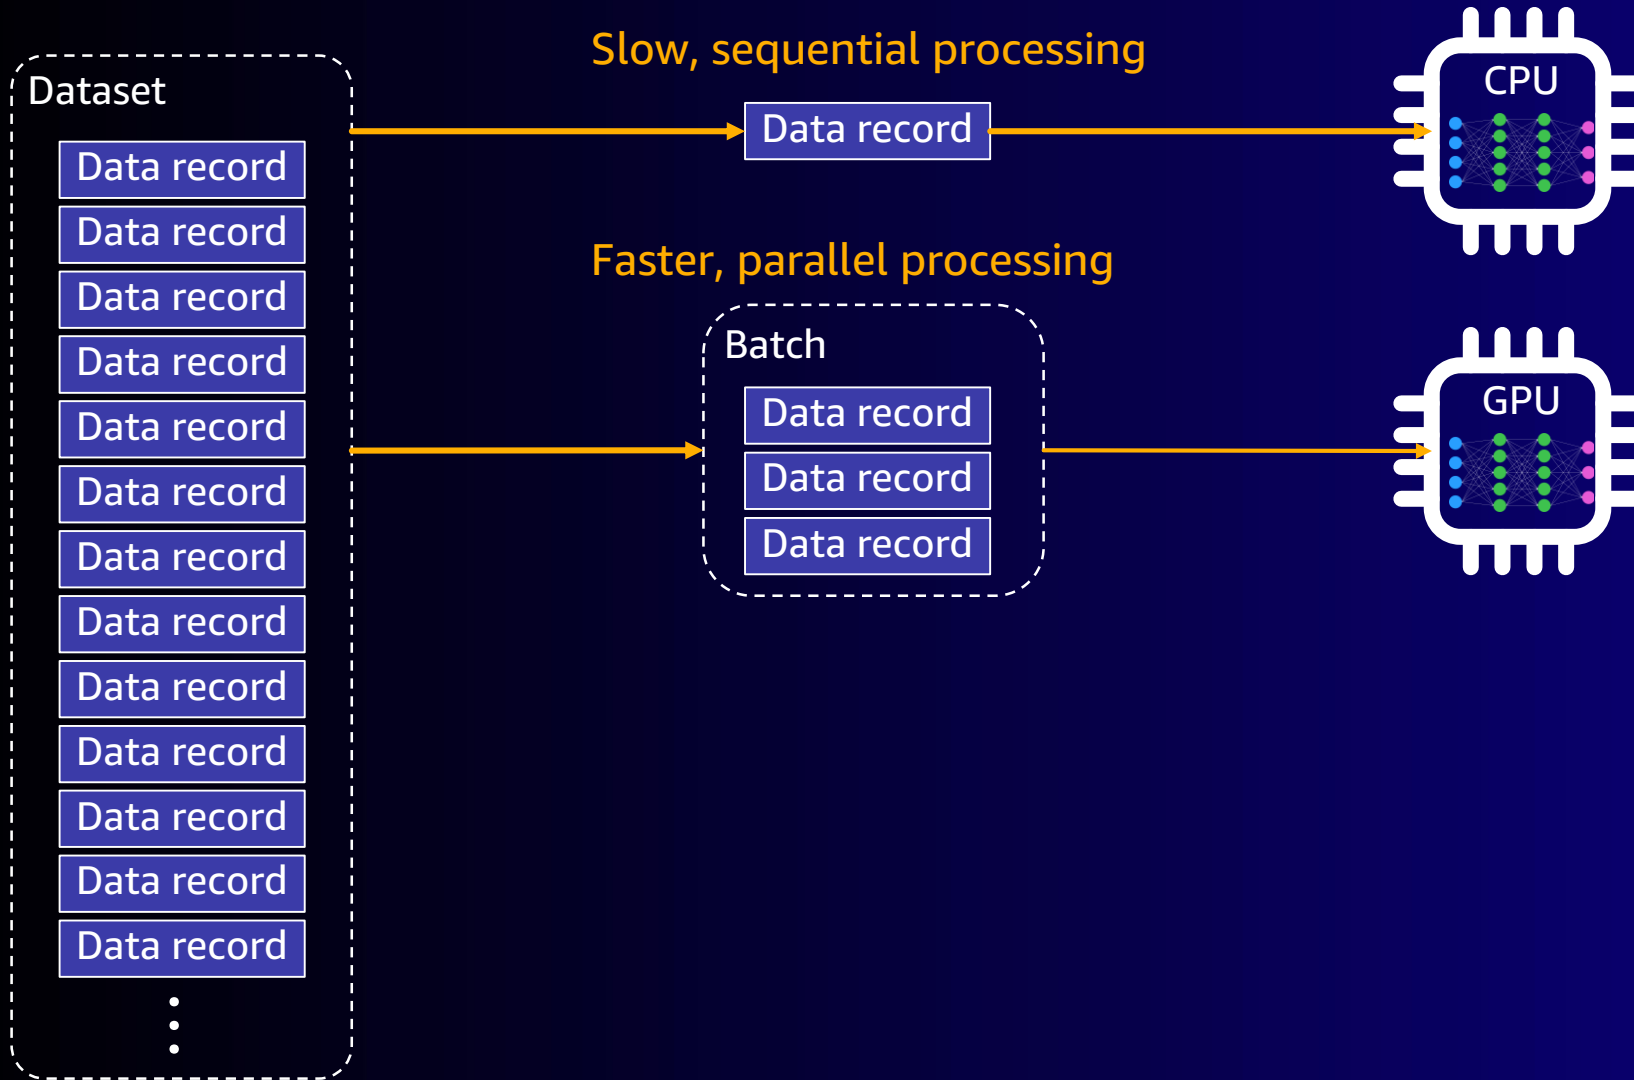
The use case for data parallelism



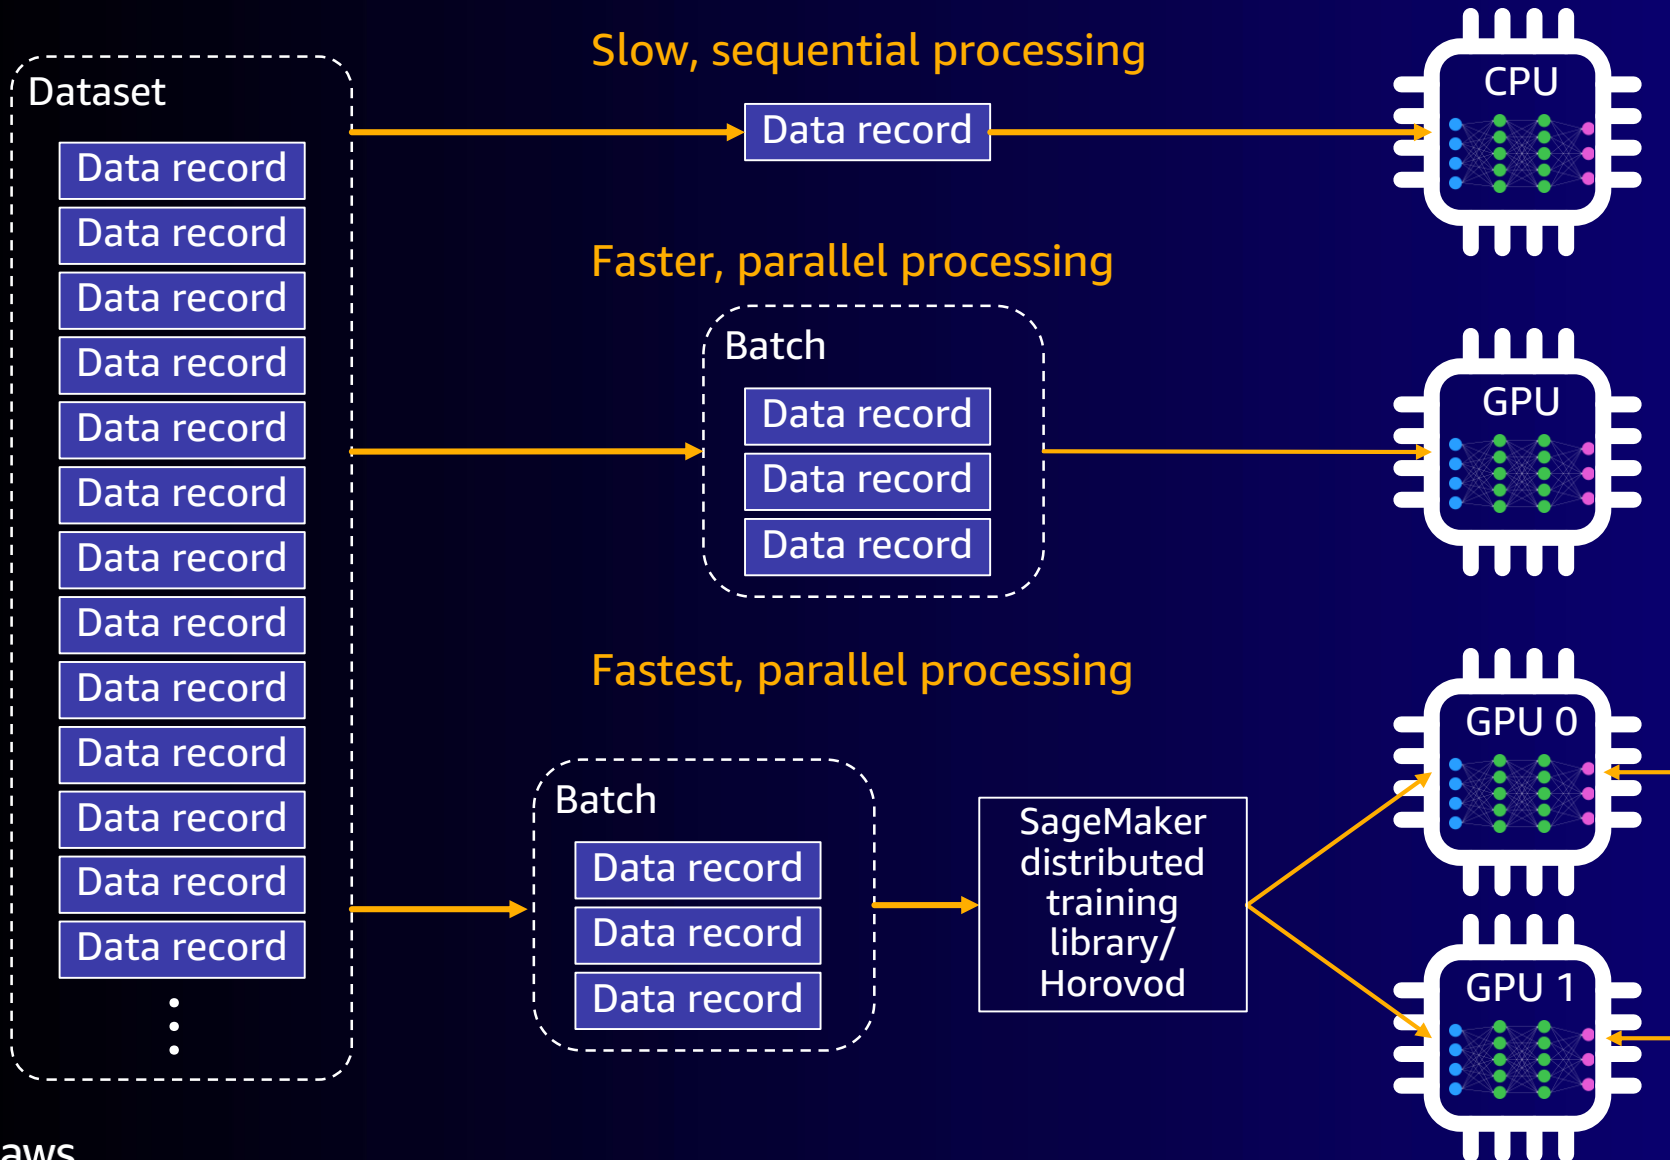
The use case for data parallelism



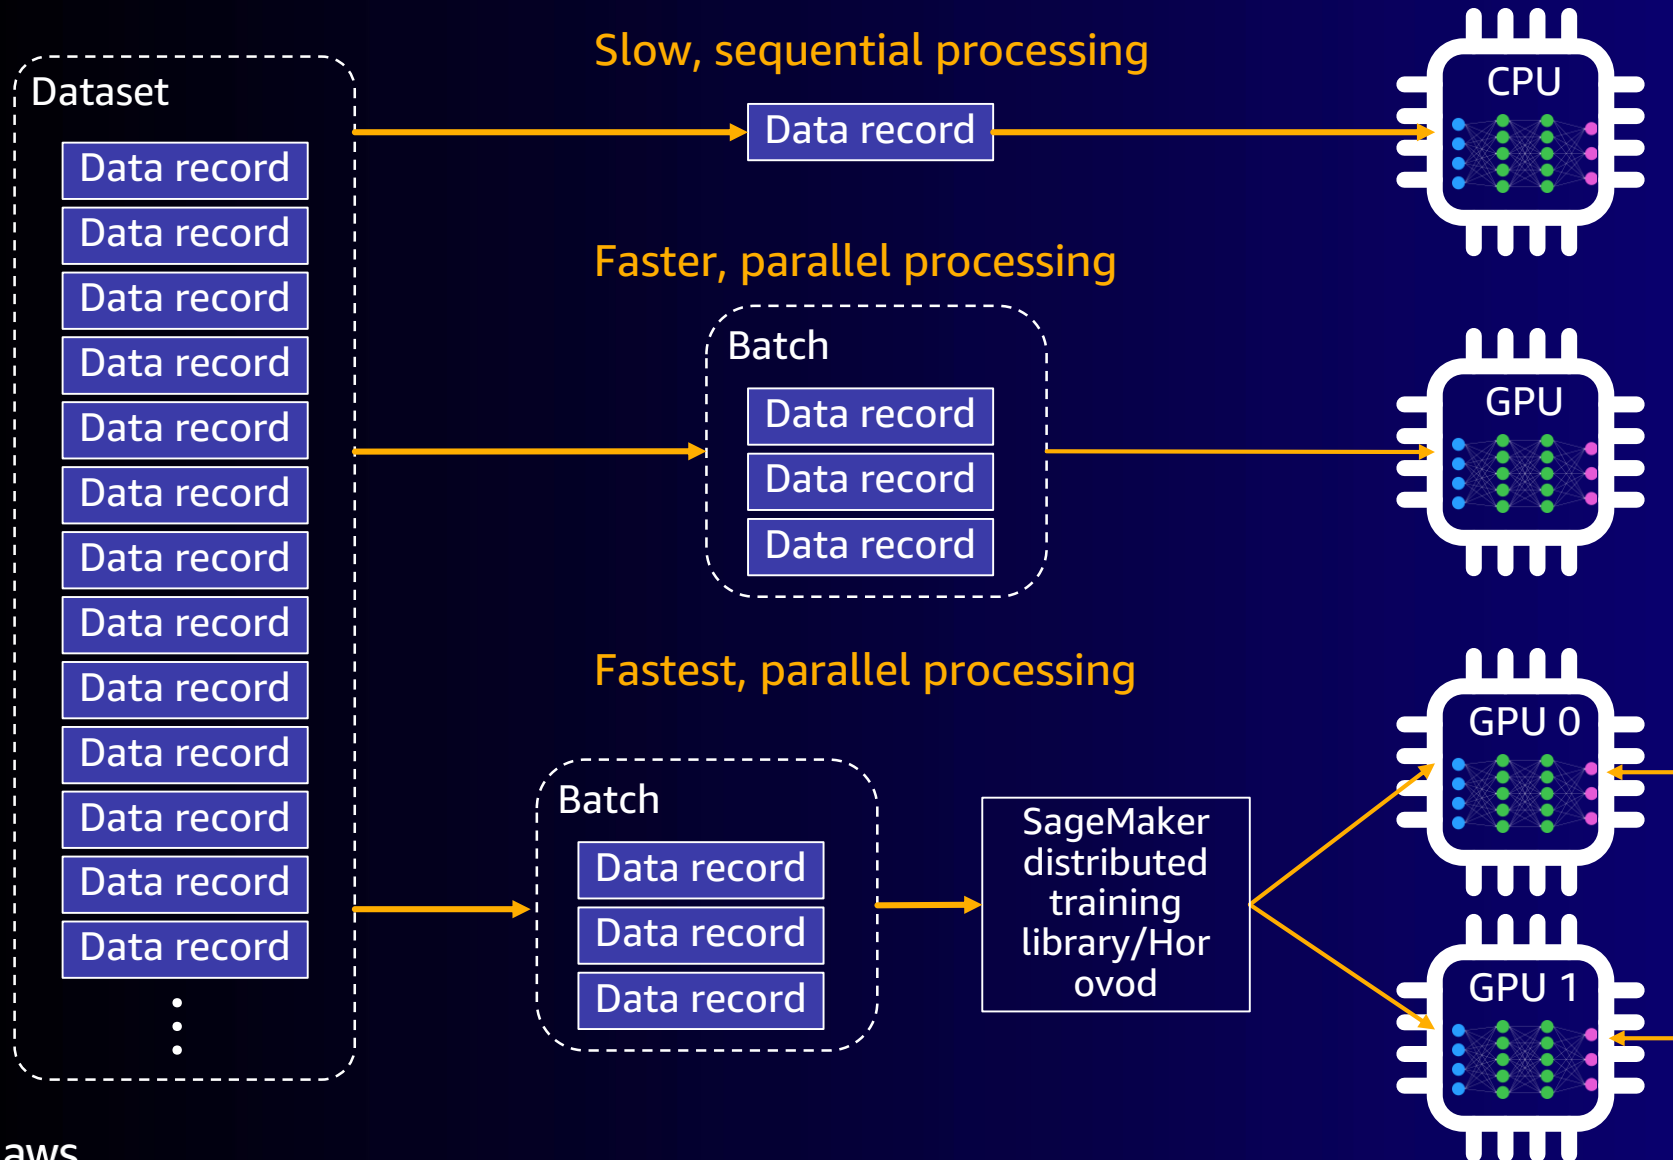
The use case for data parallelism



The use case for data parallelism



The use case for data parallelism



Combine with:



Amazon Elastic
File System



Amazon FSx
for Lustre



Amazon S3

About Horovod

- Uber's open-source framework for distributed deep learning
- Available for use with most popular deep learning toolkits like TensorFlow, Keras, PyTorch, and Apache MXNet
- Uses the all-reduce algorithm for fast distributed training rather than using a parameter server approach
- Includes multiple optimization methods to make distributed training faster

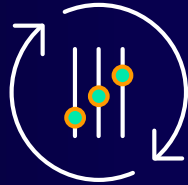
Storage options with Amazon SageMaker Training

Amazon FSx for Lustre



Amazon FSx
for Lustre

High-performance
workloads



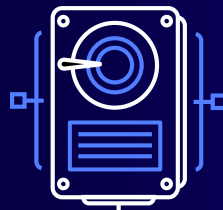
Fully
managed



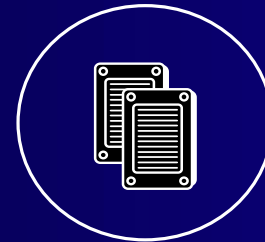
Very high
performance



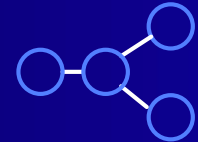
Amazon S3 data
sets as POSIX
file system



Cost
optimized



Scratch or
Persistent
mode

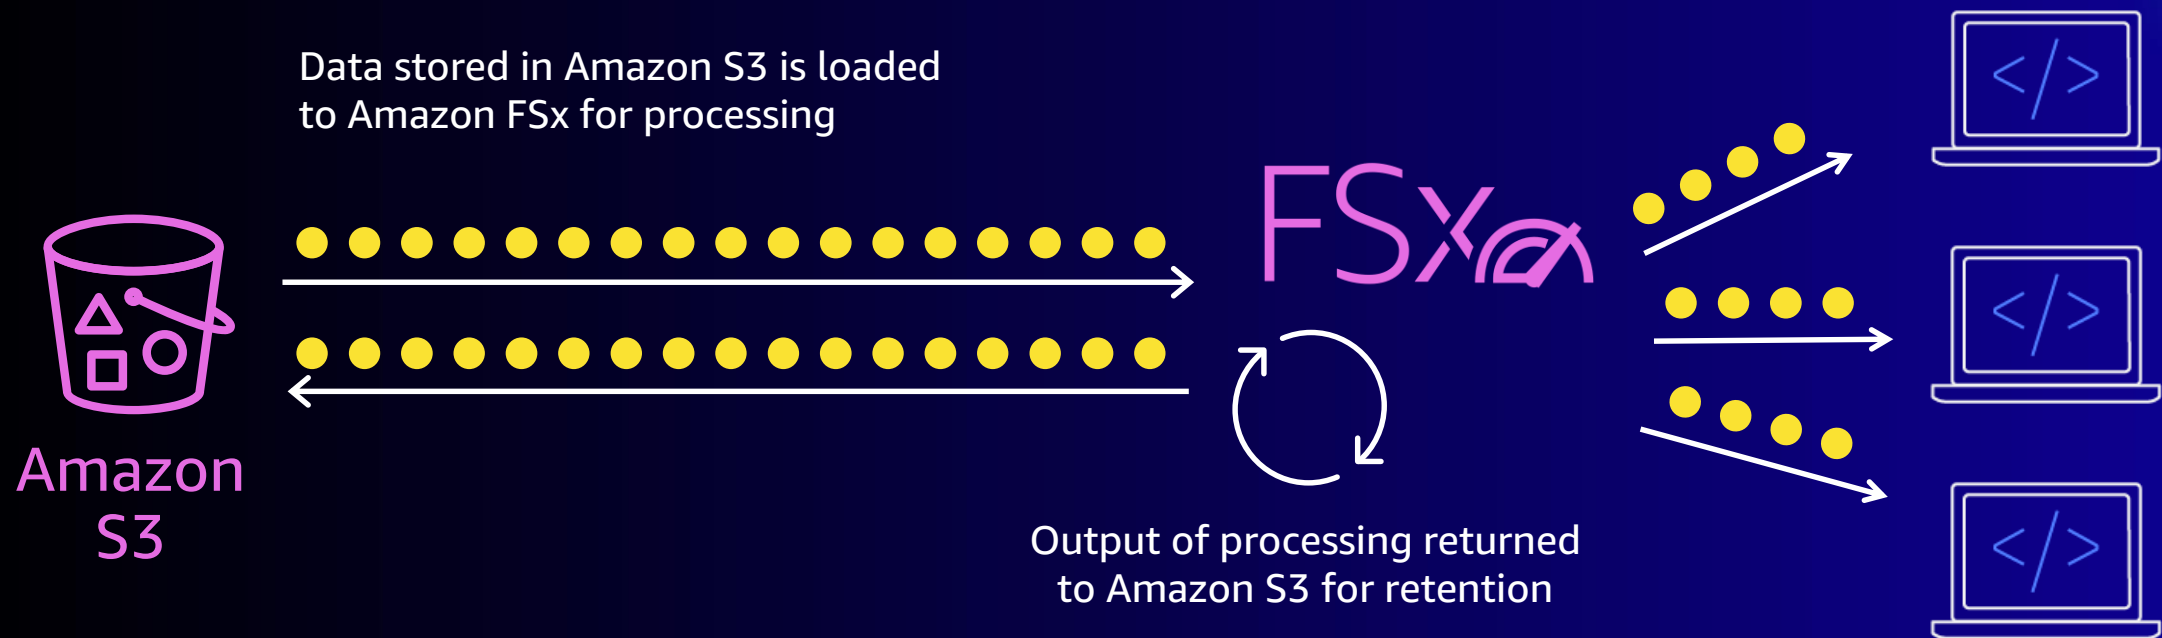


Access from
on-premises &
AWS services

Seamless integration with Amazon S3

Link your Amazon S3 dataset to your FSx for Lustre file system, and then

...



When your workload finishes, simply delete your file system

What is Amazon Elastic File System?

SIMPLE, SCALABLE, FULLY MANAGED, HIGHLY DURABLE, AND AVAILABLE SHARED FILE SYSTEM FOR AWS COMPUTE

Simple and highly reliable



Elastic

Pay only for capacity used
Performance built in – scales with capacity

Serverless shared storage



Serverless and scalable

No provisioning, scale capacity,
connections, and IOPS

Performant and cost optimized



Performant

10s of GB/s of and 500,000+ IOPS



3-AZ durable and all-AZ available

Designed for 99.999999999% durability
99.99% availability SLA



Concurrent access for tens of thousands of connections

Amazon EC2 instances, containers,
and AWS Lambda invocations



Two storage classes

Lifecycle-based cost optimization

Amazon S3



- **Exabytes** of storage
- Holds more than **100 trillion objects**
- **Tens of thousands** of data lakes built on Amazon S3
- Regularly peak at **tens of millions of requests** per second
- S3 Intelligent-Tiering is the only storage class that **automatically saves up to 95% on storage costs**
- S3 Glacier Deep Archive offers the **lowest-cost cloud object storage**, at \$0.00099 per GB-month

Amazon SageMaker with different storage options: Demo

Demo details - distributed TensorFlow training for **Mask-RCNN**



Observations from demo

Demo - Observations

Storage type	Example applications
Amazon S3	Each time the SageMaker training job is launched, it takes approximately 20 minutes to download 25 GB of training dataset from our S3 bucket to the <i>Amazon EBS volume</i> attached to each training instance. During training, data is input to the training data pipeline from the EBS volume attached to each training instance.
Amazon EFS	It takes approximately 46 minutes to copy training dataset from our S3 bucket to our EFS file-system. We only need to copy this data once. During training, data is input from the shared <i>Amazon EFS file-system</i> mounted on all the training instances.
Amazon FSx Lustre	It takes approximately 10 minutes to create a new FSx Lustre file-system and import training dataset from our S3 bucket to the new FSx Lustre file-system. You only need to do this once. During training, data is input from the shared <i>Amazon FSx Lustre file-system</i> mounted on all the training instances.

Amazon SageMaker Training – Storage Options – S3

FILE MODE - COPY DATASET FROM S3 TO INSTANCE DISK

Pros

- Simple to use
- No brainer for small datasets < 50GB
- Compatible with SageMaker local mode
- Exposes R/W POSIX filesystem
- Great for both sequential or random reads
- Benefits from filesystem caching
- Supports S3prefix, manifest file, augmented manifest file

Cons

- Idle time as training starts after all data channels are copied
- Pretty slow (~170-200MB/sec – 100GB/10min)
- Not adequate for very large datasets (>100GB) as instance is idle while copying data (spot interruption requires re-copy)

Amazon SageMaker Training – Storage Options – File

AMAZON EFS – SAGEMAKER MOUNTS THE EXISTING FILESYSTEM ON THE INSTANCE

Pros

- Great if your datasource lives in EFS rather than in S3
- Exposes a POSIX filesystem

Cons

- Not simple to setup (FS creation, VPC, security groups, etc)

Amazon SageMaker Training – Storage Options – FSx

AMAZON FSX LUSTRE – SAGEMAKER MOUNTS THE EXISTING FILESYSTEM ON THE INSTANCE

Pros

- Built for HPC: Can scale out as needed
- Great for large datasets
- Compatible with SageMaker local mode
- Exposes a POSIX filesystem
- Can do random reads fairly well
- Lazy loaded from S3, dataset can be larger than the filesystem size

Cons

- Not simple to setup (FS creation, VPC, security groups, etc)
- Cost – can be pretty expensive to maintain 24x7 (\$140/TB/Month), on top of the S3 cost.
- Need to setup/tier down before/after training if not running 24x7

Recap

Recap / other resources

1. Learn about Amazon FSx for Lustre:
<https://aws.amazon.com/fsx/lustre>
2. Learn about Amazon distributed training on Amazon SageMaker:
<https://aws.amazon.com/sagemaker/distributed-training/>
3. Mask-RCN and other Amazon SageMaker examples:
<https://github.com/aws/amazon-sagemaker-examples>

Visit the AI & Machine Learning resource hub for more resources

Dive deeper into these resources, get inspired and learn how you can use AI and machine learning to accelerate your business outcomes.

- The machine learning journey e-book
- 7 leading machine learning use cases e-book
- A strategic playbook for data, analytics, and machine learning e-book
- Accelerate machine learning innovation with the right cloud services & infrastructure e-book
- Choosing the right compute infrastructure for machine learning e-book
- Improving service and reducing costs in contact centers e-book
- Why ML is essential in your fight against online fraud e-book
- ... and more!



<https://bit.ly/3mwi59V>

Visit resource hub

AWS Machine Learning (ML) Training and Certification



AWS is how you build machine learning skills

Courses built on the curriculum leveraged by Amazon's own teams. Learn from the experts at AWS.

aws.training/machinelearning



Flexibility to learn your way

Learn online with on-demand digital courses or live with virtual instructor-led training, plus hands-on labs and opportunities for practical application.

explore.skillbuilder.aws/learn



Validate your expertise

Demonstrate expertise in building, training, tuning, and deploying machine learning models with an industry-recognized credential.

aws.amazon.com/certification

Thank you for attending AWS Innovate – AI/ML 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!

Gaurav Singh

