

The logo for AWS Summit, featuring the lowercase 'aws' with the Amazon smile arrow, followed by the word 'SUMMIT' in all caps.

# aws SUMMIT

AIM309

# Train and tune foundation models at scale with Amazon SageMaker

**Kanwaljit Khurmi**

Principal GenAI/ML Solutions Architect  
AWS

**Sean Smith**

Sr. WW Specialist SA GenAI  
AWS




# Customer segments served by AWS

Strong control and flexibility


Low entry cost and complexity, faster TTM

## Building/retraining FMs from scratch



Build FMs or retrain open source FMs from scratch

## Finetuning FMs



Finetune FMs for specific domain/workload

## Customizing and Using FMs

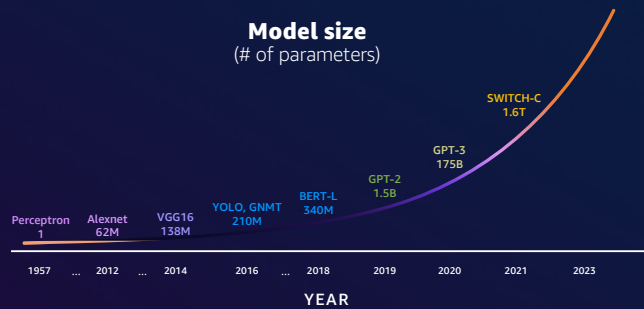


Enhance finetuned FMs with context and use them


# Our Observation

Develop and train continuously for weeks and months on SageMaker-managed infrastructure

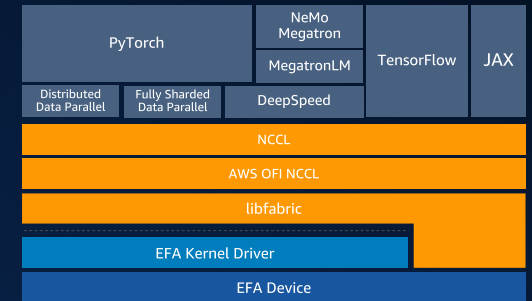
## Performance (time & cost)



## Resilience



## Ease of use & flexibility




# Amazon SageMaker Training



## Resilient environment

---

Self-healing clusters reduce training time by up to 20%



## Streamline distributed training

---

SageMaker distributed training libraries improve performance by up to 20%



## Purpose built ML infrastructure

---

A broad choice of GPUs and CPUs, as well as AWS accelerators such as AWS Trainium and AWS Inferentia

# Amazon SageMaker for large scale FM Training:

## Create training jobs on SageMaker

---

Focus on ML rather than IT to accelerate time to market

Access to flexible on-demand GPU cluster

## Build on Amazon SageMaker HyperPod

---

Customize and manage cluster orchestration using Slurm

# Amazon SageMaker for Large Scale FM Training:



## Resilient environment

---

Self-healing clusters reduce training time by up to 20%

AND



## Streamline distributed training

---

SageMaker distributed training libraries improve performance by up to 20%

AND



## Managed training environment (SMTJ)

---

Focus on ML without the need to manage infrastructure

OR

## Optimized resources utilization (SMHP)

---

Control over compute environment and workload scheduling

# Foundation model training on Amazon SageMaker

## OPTIMIZED DISTRIBUTED TRAINING LIBRARIES & FRAMEWORKS



TensorFlow



PyTorch



Hugging Face

SageMaker Distributed  
Training Libraries

Bring your own library (e.g.,  
Nemo Megatron, FSDP)

## AMAZON SAGEMAKER TRAINING

Large-scale cluster  
orchestration

NCCL health checks

Resilient training

On-demand and  
persistent clusters

Warm pools

SSH to container

Efficient data loading

Debugger

Profiling

Experiment tracking

Hyperparameter  
optimization

Pay for what  
you use

## ML COMPUTE INSTANCES & ACCELERATORS

NVIDIA GPUS  
H100, A100, A10

AWS Trainium  
Trn1 and Trn1n

CPU instances

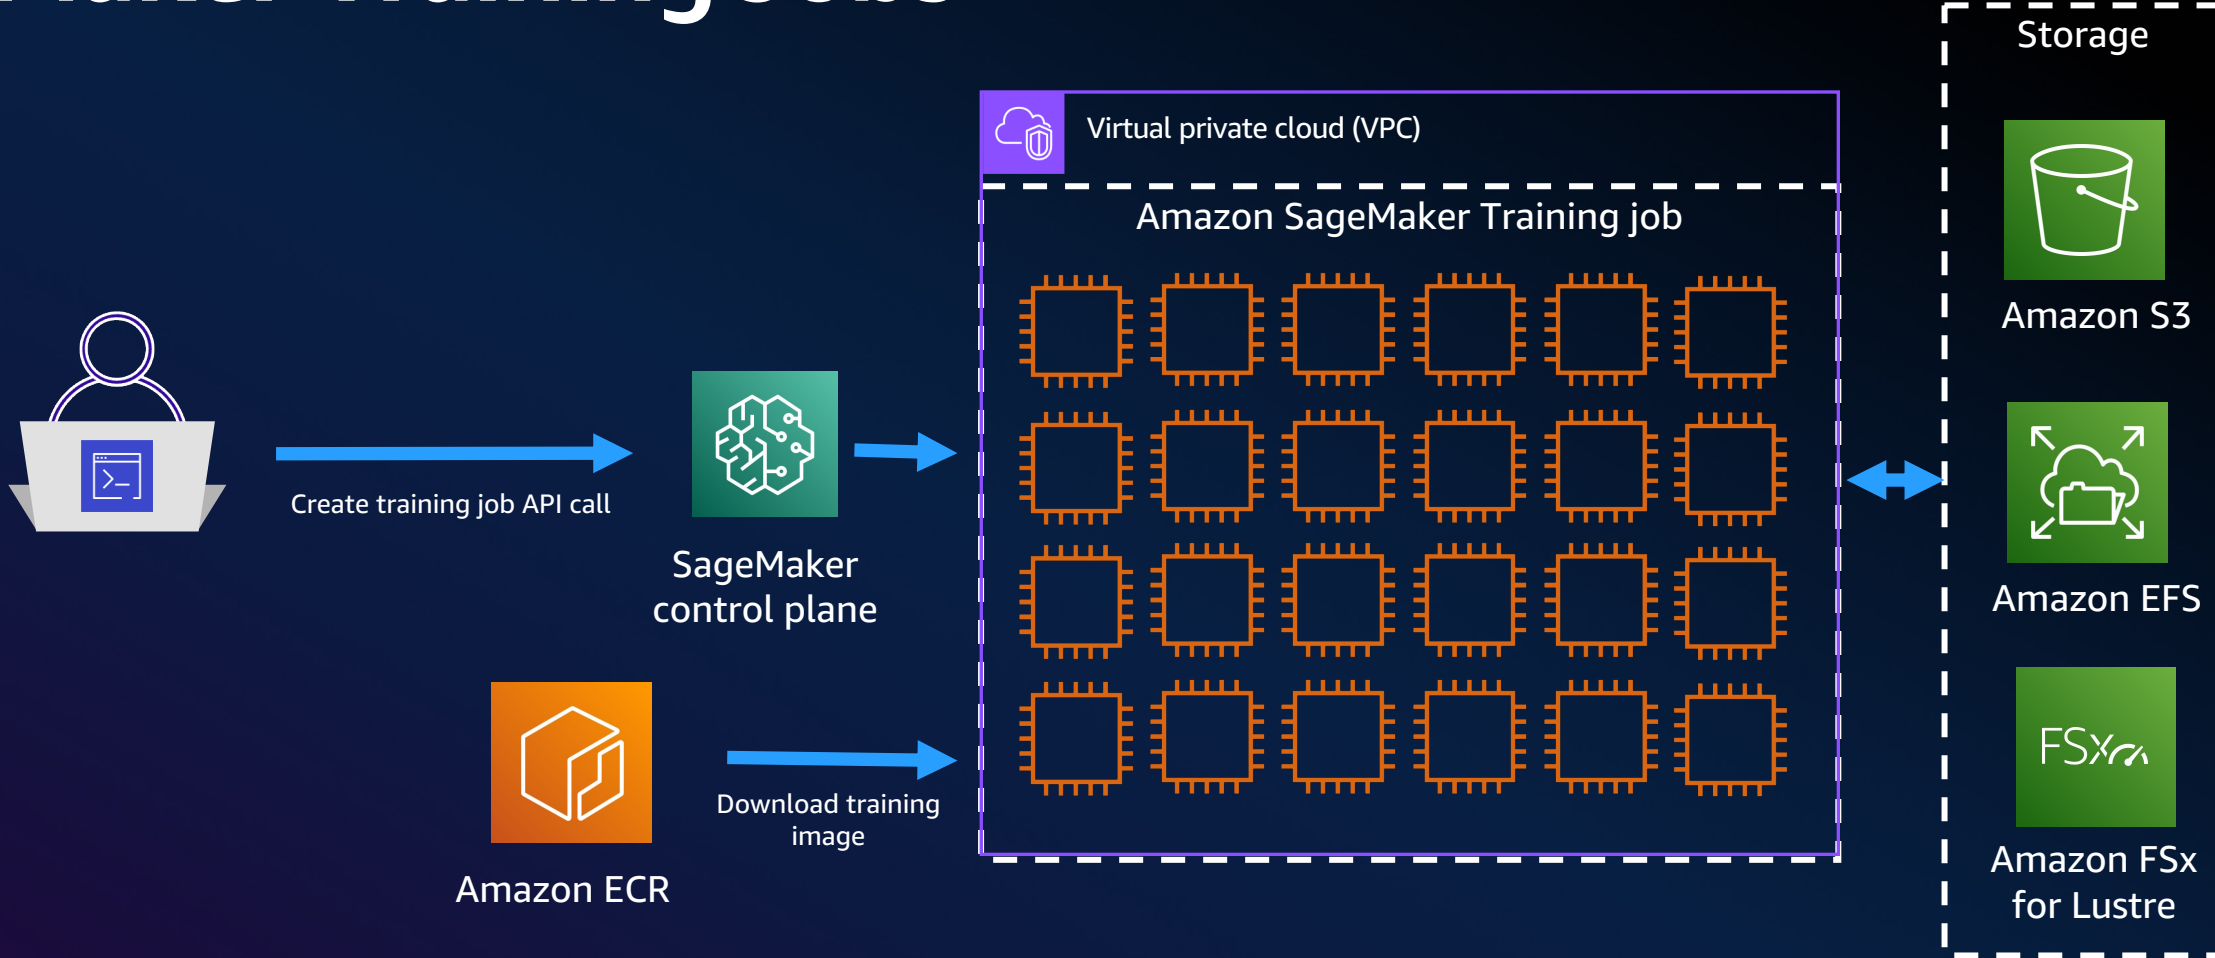
AWS Nitro System

EFA based networking





# SageMaker Training Jobs



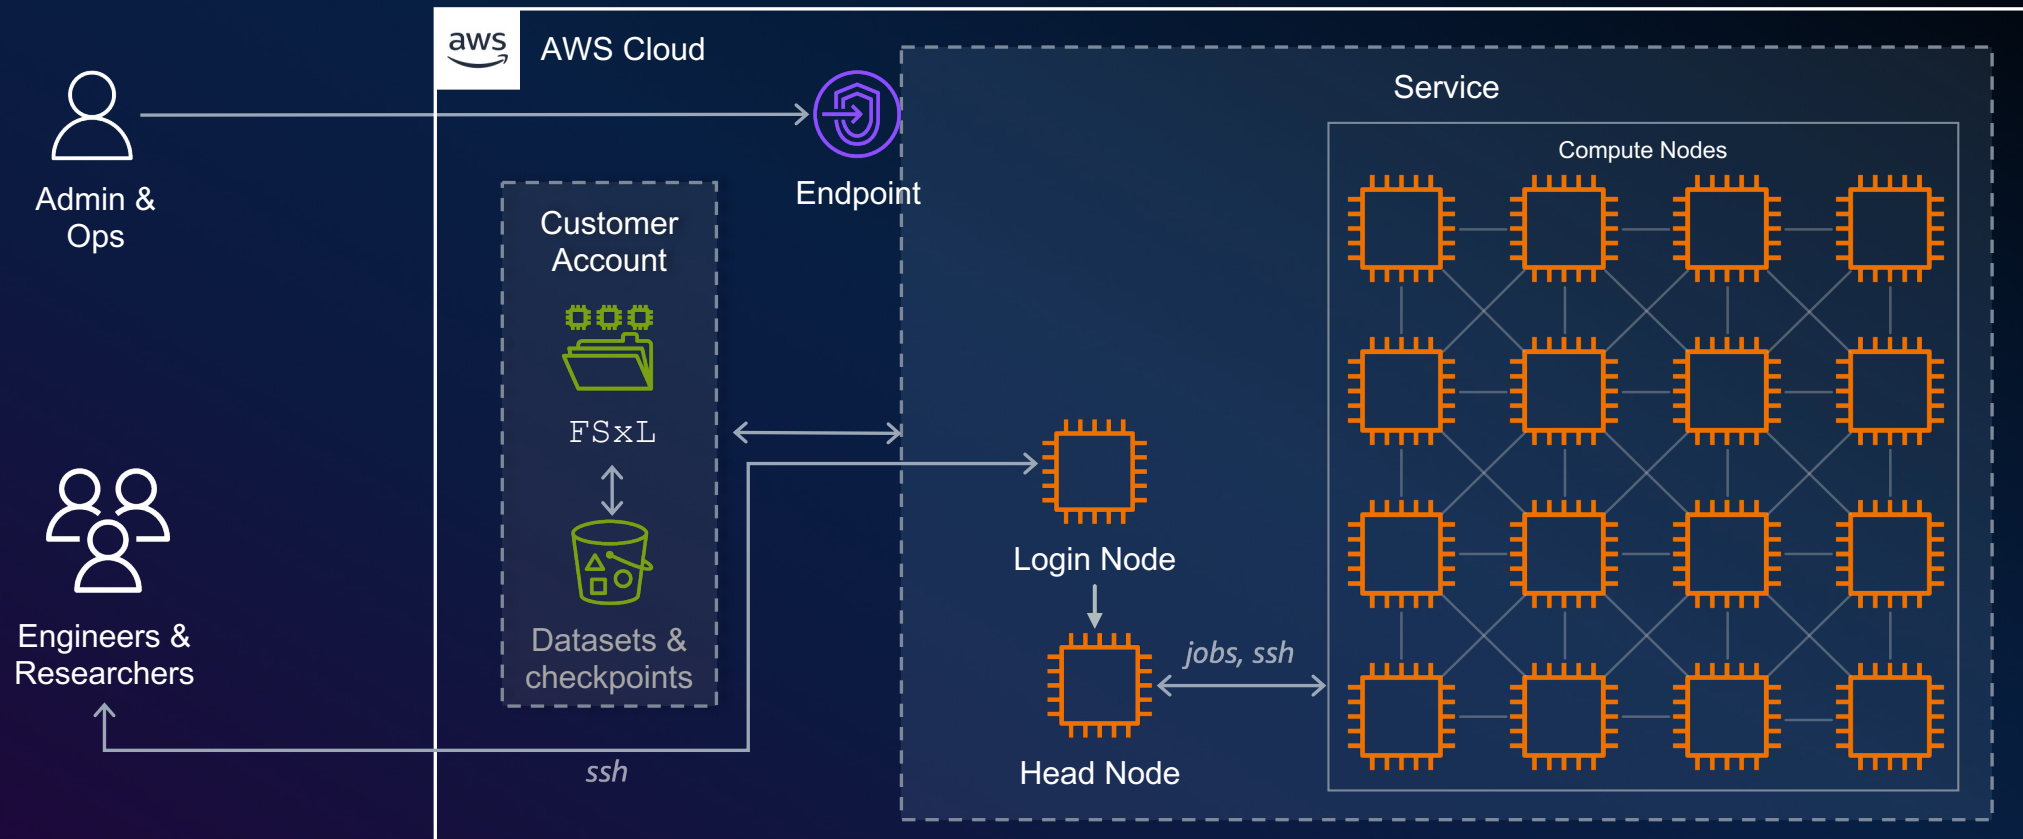
- **Managed training environment**
- **Resilient environment**
- **Streamline distributed training**



# SageMaker Training Jobs Demo



# SageMaker HyperPod example architecture



# SageMaker HyperPod Demo



# Thank you!

**Kanwaljit Khurmi**

kkhurmi@amazon.com

**Sean Smith**

seam@amazon.com



Please complete the session survey in the mobile app

# SageMaker Training Jobs:

Proven, large Scale, resilient training clusters - for high-volume training jobs with 100s of billions of parameters

[Customer Stories](#) / Software & Internet



2023



## AI21 Labs Accelerates Generative AI Model Adoption Using Amazon SageMaker

Learn how AI21 Labs, a leader in generative AI and large language models, rapidly pretrained and released a 17-billion-parameter model using Amazon SageMaker.

[Overview](#) | [Opportunity](#) | [Solution](#) | [Outcome](#) | [AWS Services Used](#)



[AWS Machine Learning Blog](#)

## Technology Innovation Institute trains the state-of-the-art Falcon LLM 40B foundation model on Amazon SageMaker

by Dr. Ebtesam Almazrouei, Olivier Cruchant, and Will Badr | on 07 JUN 2023 | in [Amazon SageMaker](#), [Artificial Intelligence](#), [Intermediate](#) (200) | [Permalink](#) | [Comments](#) | [Share](#)

*This blog post is co-written with Dr. Ebtesam Almazrouei, Executive Director–Acting Chief AI Researcher of the AI-Cross Center Unit and Project Lead for LLM Projects at TII.*

United Arab Emirates’s (UAE) [Technology Innovation Institute \(TII\)](#), the applied research pillar of Abu Dhabi’s [Advanced Technology Research Council](#), has launched Falcon LLM, a foundational large language model (LLM) with 40 billion parameters. TII is a leading global research center dedicated to pushing the frontiers of knowledge. TII’s team of scientists, researchers, and engineers work to deliver discovery science and transformative technologies. TII’s work focuses on breakthroughs that will future-proof our society. Trained on 1 trillion tokens, [TII Falcon LLM](#) boasts top-notch performance while remaining incredibly cost-effective. **Falcon-40B matches the performance of other high-performing LLMs, and is the top-ranked open-source model in the public [Hugging Face Open LLM leaderboard](#).** It’s available as open-source in two different sizes – Falcon-40B and Falcon-7B and was built from scratch using data preprocessing and model training jobs built on [Amazon SageMaker](#). Open-sourcing Falcon 40B enables users to construct and customize AI

[Customer Stories](#) / Research and Development


2023



## LG AI Research Develops Foundation Model Using Amazon SageMaker

LG AI Research built EXAONE—a foundation model that can be used to transform business processes—using Amazon SageMaker, broadening access to AI in various industries such as fashion, manufacturing, research, education, and finance.

[Overview](#) | [Opportunity](#) | [Solution](#) | [Outcome](#) | [AWS Services Used](#) | [Architecture Diagram](#)



March 21, 2024

## Upstage Launches Small Language Model on AWS to Help Businesses Around the World Build and Scale Generative AI Applications for the Korean Market

*[MEDIA ALERT]*

*Leading South Korean AI company’s language model is highly proficient and culturally aware, accelerating the development of use cases across a broad range of industries including entertainment, healthcare, research, ecommerce, and education*

*The model is available today in Korean and English (Thai and Japanese to be available within the year) on Amazon JumpStart and AWS Marketplace*





# Resilient computing with SageMaker HyperPod




## Stability AI

- Leading open-source GenAI company
- **Challenge:** Losing hundreds of hours to hardware failures



HyperPod automatically replaces faulty instances once detected



Workloads are automatically resumed – saving Stability AI over 50% in training time and cost



# Performant computing with SageMaker HyperPod



## Perplexity AI

- First conversational Q&A engine
- **Challenge:** low throughput of LLM training experiments



SageMaker distributed training libraries provide an optimized backend for deep learning models in TensorFlow, PyTorch



Up to 20% faster and cheaper than NCCL and MPI-based solutions – **Unlocking a 2x increase in experiment throughput for Perplexity AI**






# Customizable computing with SageMaker HyperPod




## Hugging Face

- Open source machine learning platform and community
- **Challenge:** Research requires control over all layers of tech stack



Customize HyperPod clusters with frameworks, containers, debugging tools



Hugging Face's science teams has the **freedom to innovate quickly**