



UNLOCK YOUR

# AI ADVANTAGE

RWS + AWS

Mark Lawyer



# RWS Clients

88

of the globe's  
top 100 brands

8,000+  
active enterprise clients

## Sector expertise

Automotive  
Aerospace  
Chemical  
Defence  
Financial  
Government  
Legal  
Manufacturing  
Medical devices  
Pharmaceutical  
Retail  
Technology  
Telecommunications  
Travel & Tourism

47 AI-  
related  
patents

16 of  
Fortune's  
top 20  
most  
admired  
companies

19 of the  
top 20  
pharma  
companies

Support for  
548  
different  
language  
pairs

Translating  
nearly 3  
billion  
words a  
year

18 of the  
top 20 law  
firms

# AI

Threat or  
opportunity?



RWS

A close-up photograph of a person's face, smiling and wearing black-rimmed glasses. The image is overlaid with a semi-transparent dark blue and purple gradient. The letters 'AI' are prominently displayed in white, bold, sans-serif font over the person's eyes and nose.

# AI

Content and  
Language are top  
of the agenda in board  
rooms around the world

Overwhelming **GROWTH**  
**IN CONTENT**

Volume and Complexity





How to **manage it, deliver it and monitor it** to

**ACHIEVE LOCAL  
IMPACT AND TRUST  
on a  
GLOBAL SCALE**

# RWS+AWS

## Innovation Model





**Amazon SageMaker**



Language Weaver

Foundation  
models



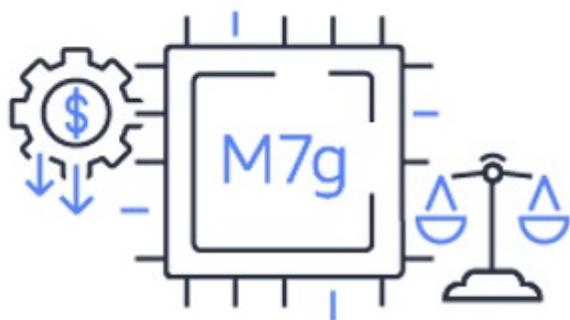
**Amazon Bedrock**



Trados

Custom model  
training

Scalability



**Amazon Graviton**



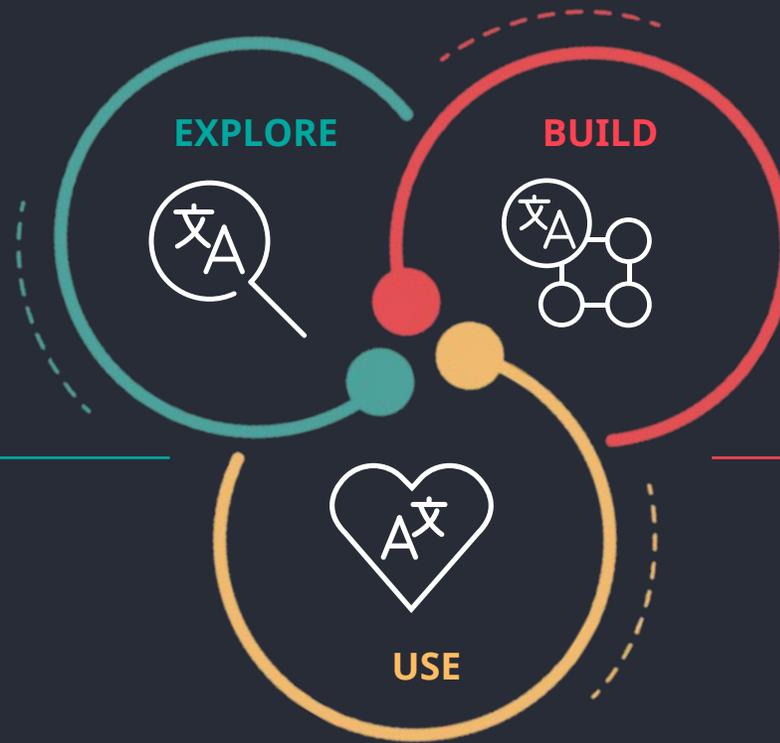
Tridion

Security

# INNOVATION TO UNLOCK YOUR AI POTENTIAL

Choose the right AI strategies and tools

TECH SERVICES



Train your AI with dependable, responsible data

TRAIN AI

Build a secure linguistic AI platform tailored for your business

LANGUAGE WEAVER

Understand content in any language, instantly

LANGUAGE WEAVER

Extend your use of AI beyond machine translation and content analysis - improving productivity and automation

LINGUISTIC SERVICES

Use our end-to-end translation management platform to deliver translation projects faster and smarter

TRADOS

Author, manage, collaborate and publish your content in our CCMS

TRIDION + FONTO + PROPYLON



powered by aws

**AI = OPPORTUNITY**



# Keynote : Thomas' take (25 min)

- What does this Partnership Mean from a RWS prospective?
- Future Innovation with AWS as a Tech Partner.
- Strategic Collaboration agreement
- Launch of Evolve with an AWS backbone
- Other Product innovations like Tridion Docs with AI
- Launch any POC programmes
- Joint GTM with AWS.

*Thank you*

BACK AGAIN



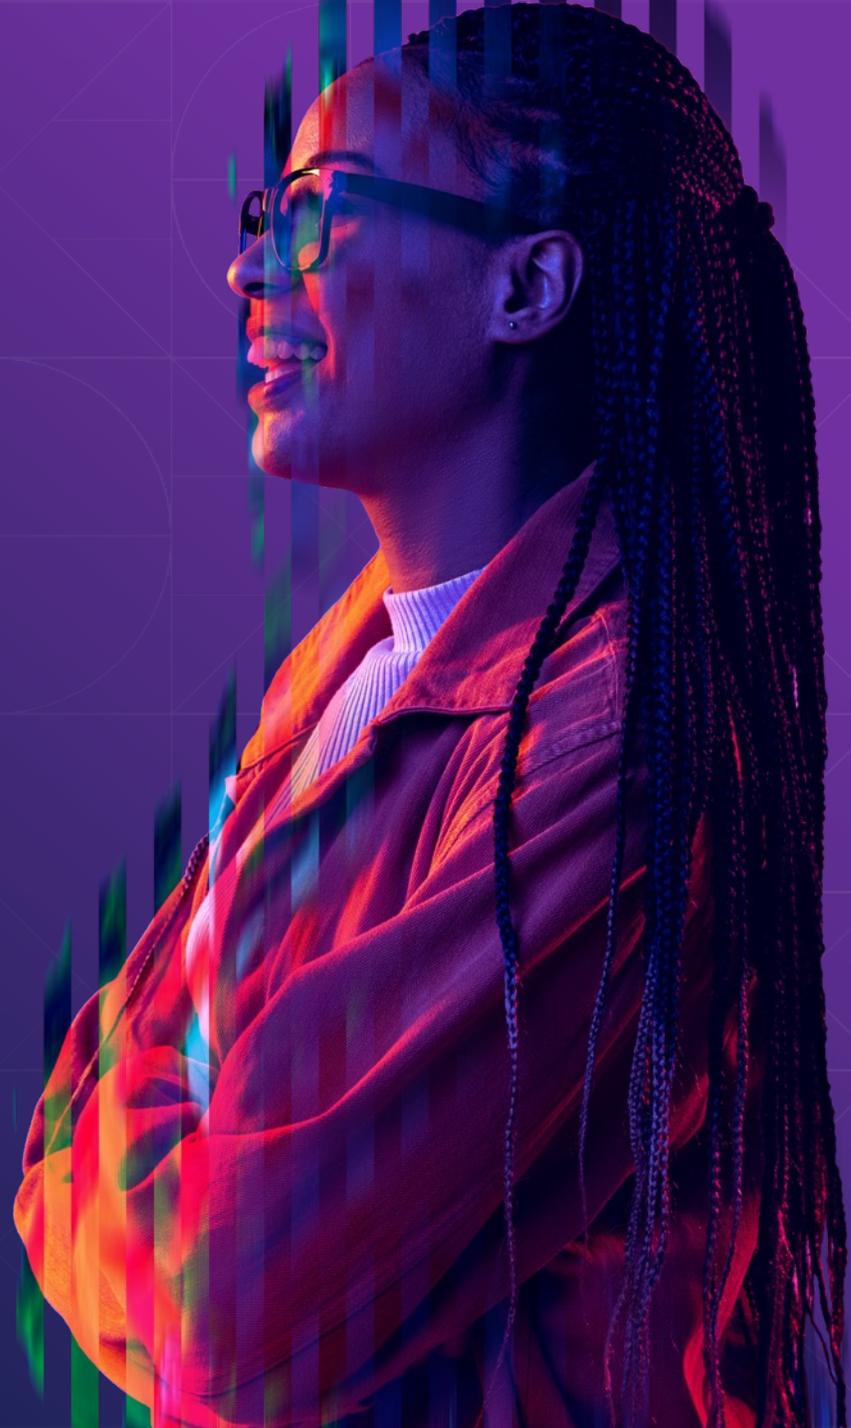
# RWS Overview

<b>OPERATING DIVISIONS</b>	<b>Language Services</b> <ul style="list-style-type: none"><li>• Localisation solutions to multiple verticals</li><li>• Includes TrainAI data services, eLearning, video localisation and interpreting services</li></ul>	<b>Regulated Industries</b> <ul style="list-style-type: none"><li>• Life sciences</li><li>• Financial services</li><li>• Legal services</li><li>• Highly specialised technical translations</li></ul>	<b>Language &amp; Content Technology</b> <ul style="list-style-type: none"><li>• Linguistic AI<ul style="list-style-type: none"><li>. Language Weaver</li><li>. Trados</li></ul></li><li>• Content Management<ul style="list-style-type: none"><li>Tridion, Fonto, Propylon, Contenta</li></ul></li></ul>	<b>IP Services</b> <ul style="list-style-type: none"><li>• Patent translation and filing</li><li>• Search, retrieval and monitoring services</li><li>• Highly specialised technical translations</li></ul>
<b>REVENUE SPLIT</b>	<b>HY23: 44%</b>	<b>HY23: 23%</b>	<b>HY23: 18%</b>	<b>HY23: 14%</b>

# Managing our impact

*Eric Hodge*

*VP Security, Service Delivery  
Cloud Operations*



# Innovative Engineering at RWS



## CONTENT

- Improving technical documentation
- Creating engaging digital experiences
- Optimizing content processes
- Succeeding with digital workplaces



## INNOVATION

- Neural machine translation
- Multilingual content intelligence
- Headless content delivery
- Trustable chat
- Content Readability



## LANGUAGE

- Localizing and translating content
- Improving translation quality
- Managing translation projects
- Delivering multilingual videos
- Boosting translation team productivity



powered by **aws**  

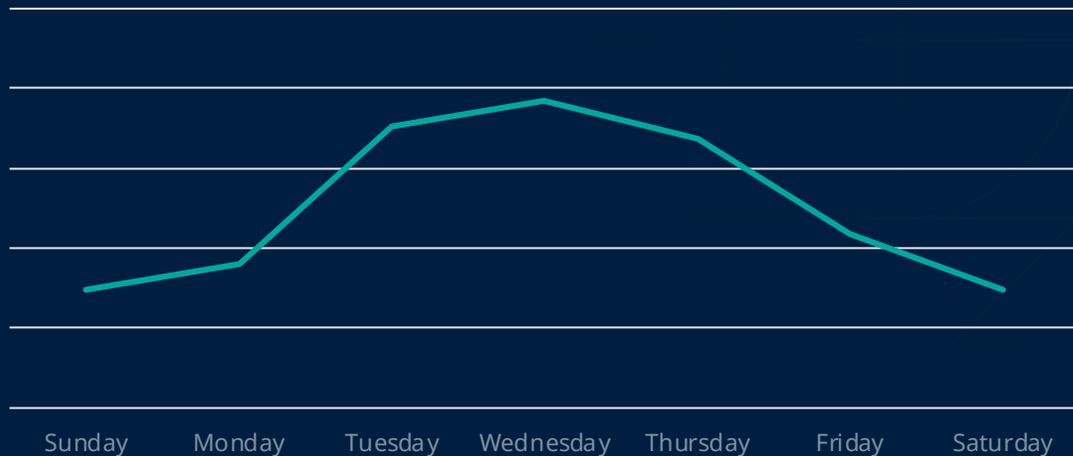



# A few numbers about RWS cloud



## Dynamic

Active CPUs



## Storage

2.3 Petabytes  
4,348 EBS volumes



## Linux Adoption

Roughly 50% running on Linux  
(from ~15% 3 years ago)



US-West-1

US-West-2

CA-Central-1

US-East-2

US-East-1

SA-East-1

EU-West-1

EU-West-2

EU-Central-1

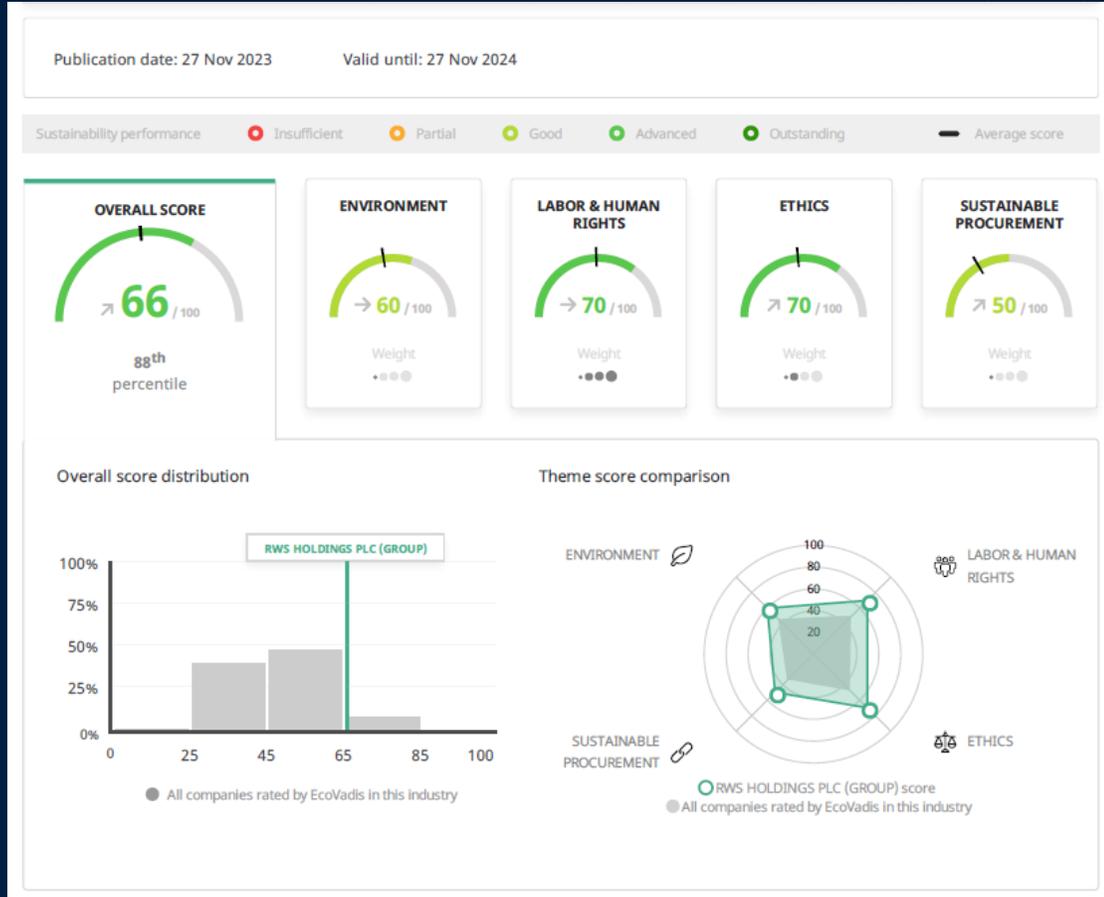
AP-South-1

AP-East-1

AP-Southeast-1

AP-Northeast-1

# ESG – Our responsibility





ALL  
YOU

NEED  
IS LESS

# As a SaaS vendor... can you really claim to be green?



## Economy of scale helps us improve daily

Faster feedback loop on our developments, increasing innovation speed.



## Riding on AWS's own targets

100% renewable energy by 2025

# How do we minimize our impact?



## Static workloads

Run the same instances for as long as we can.

Deploy latest-gen, stay there.



## Dynamic workloads

Let AWS handle it.

Use spot instances as much as possible.



## Adopt more efficient types

Deploy Graviton

Test various instance types for the different workloads

# Changing at scale – moving to ARM at RWS



## Test, test, test

POC for tooling, 3<sup>rd</sup> party products

POC for RWS own products



## Change our code

SQL Server (Standard and RDS) to Aurora

Migration from monolith Java, .NET applications to ECS, EKS, .NET core



## Operate it

Roll out, coordinate upgrades.

Adapt processes, procedures, backup plans, systems.

# The importance of being a **great** partner



## Proof of concept

Support  
Engineering help



## Collaborate

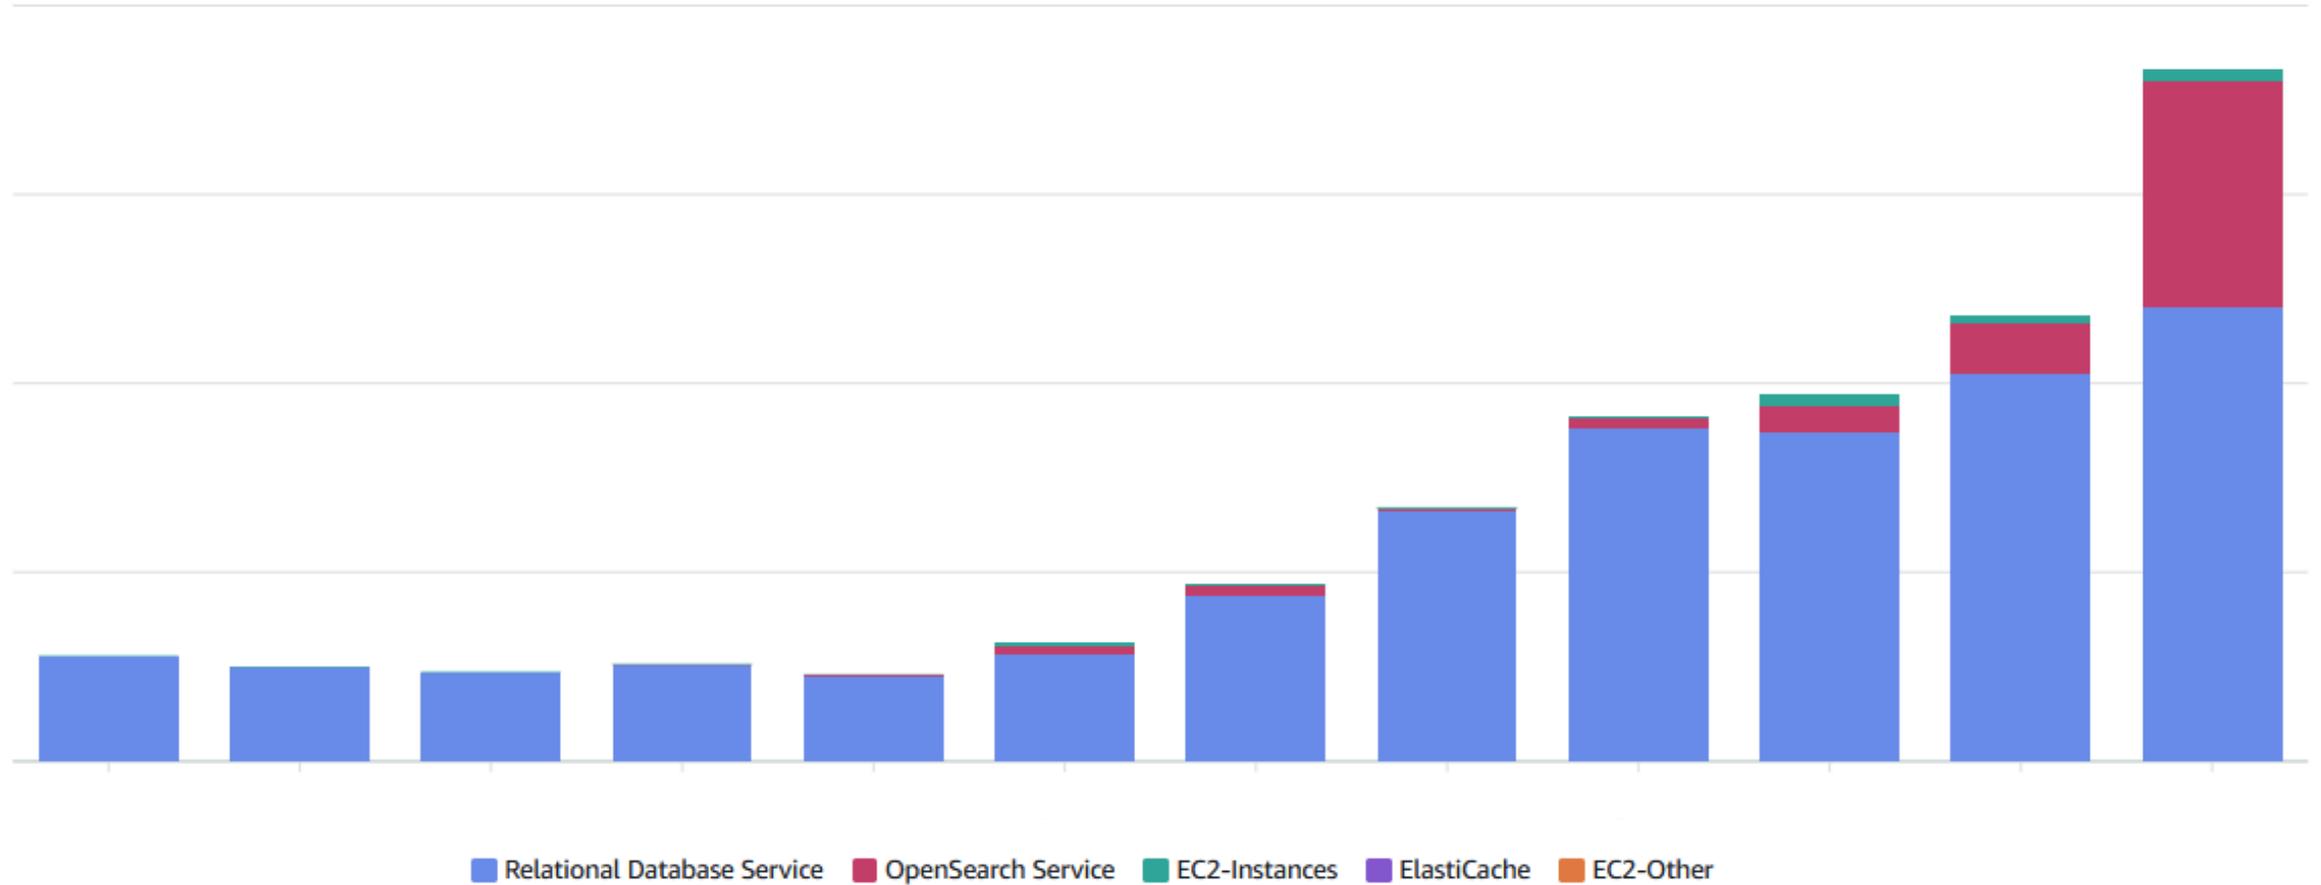
Work with R&D to address  
bottlenecks



## Measure impact

As we roll out, hold our hands,  
help us improve further

# Graviton usage across our estate - last 12 months



# Q&A



# GenAI and... Infrastructure?

Randy DeFauw

SR. PRINCIPAL SOLUTIONS ARCHITECT

AMAZON WEB SERVICES



**Enhance Customer Experiences**

**Boost employee productivity & creativity**

**Optimize business processes**

## Enhance Customer Experiences

CHATBOTS

VIRTUAL ASSISTANTS

CONVERSATION ANALYTICS

PERSONALIZATION

## Boost employee productivity & creativity

CONVERSATIONAL SEARCH

SUMMARIZATION

CONTENT CREATION

CODE GENERATION

DATA TO INSIGHTS

## Optimize business processes

DOCUMENT PROCESSING

DATA AUGMENTATION

CYBERSECURITY

PROCESS OPTIMIZATION

## Healthcare & Life Sciences

Ambient digital scribe

Medical imaging

Drug discovery

Enhance clinical trials

Research reporting

## Industrial & Manufacturing

Product design

Operational efficiency

Maintenance Assistants

Supply chain optimization

Equipment diagnostics

## Financial Services

Portfolio management

Financial documentation

Intelligent advisory

Fraud detection

Compliance assistant

## Retail

Pricing optimization

Virtual try-ons review

Marketing Optimization

Product descriptions

Pers. Recommendations

## Media & Entertainment

HQ content at scale

Enrich broadcast content

Automated content tagging

Optimize subscriber exper.

Automated highlights gen.

# Generative AI Stack



**APPLICATIONS THAT LEVERAGE LLMs AND OTHER FMs**

**TOOLS TO BUILD WITH LLMs AND OTHER FMs**

**INFRASTRUCTURE FOR FM TRAINING AND INFERENCE**



# Generative AI Stack



APPLICATIONS THAT LEVERAGE LLMs AND OTHER FMs

TOOLS TO BUILD WITH LLMs AND OTHER FMs

INFRASTRUCTURE FOR FM TRAINING AND INFERENCE





**CG1**

NVIDIA Tesla  
M2050 "Fermi"  
GPUs

**G2**

NVIDIA GRID  
GK104 "Kepler"  
GPUs

**P2**

NVIDIA  
K80  
GPUs

**G3**

NVIDIA  
Tesla M60  
GPUs

**P3**

NVIDIA V100  
Tensor Core  
GPUs

**G4**

NVIDIA T4  
Tensor Core  
GPUs

**P4**

NVIDIA A100  
Tensor Core  
GPUs

**G5**

NVIDIA A10G  
Tensor Core  
GPUs

**G5g**

NVIDIA T4G  
Tensor Core  
GPUs

**P5**

NVIDIA H100  
Tensor Core  
GPUs

A short detour into  
infrastructure

# Innovation enabled by AWS Nitro System

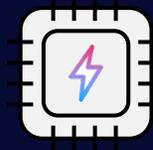
MODULAR BUILDING BLOCKS FOR RAPID DESIGN AND DELIVERY OF AMAZON EC2 INSTANCES

## NITRO CARD



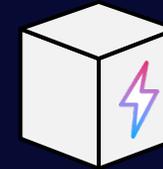
Local NVMe storage,  
Amazon EBS, networking,  
monitoring, and security

## NITRO SECURITY CHIP



Integrated into motherboard  
Protects hardware resources

## NITRO HYPERVISOR



Lightweight hypervisor  
Memory and CPU allocation  
Bare metal-like performance

# Innovating with AWS Graviton2

40% BETTER PRICE PERFORMANCE FOR A BROAD RANGE OF WORKLOADS



## BEST PRICE PERFORMANCE

Delivers up to 40% better price performance over comparable x86-based instances



## EXTENSIVE ECOSYSTEM

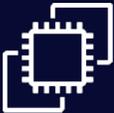
Supported by popular Linux OSes along with popular applications and services from AWS and ISVs



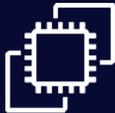
## ENHANCED SECURITY

Provide key capabilities for application security, including 256-bit DRAM encryption

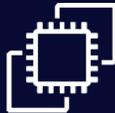
## GRAVITON2-BASED INSTANCES



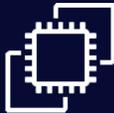
**C6G(D)**  
COMPUTE-  
OPTIMIZED



**M6G(D)**  
GENERAL  
PURPOSE



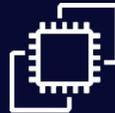
**R6G(D)**  
MEMORY-  
OPTIMIZED



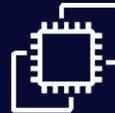
**C6GN**  
NETWORK-  
OPTIMIZED



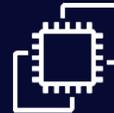
**X2GD**  
ENHANCED  
MEMORY



**T4G**  
BURSTABLE



**IM4GN**  
STORAGE-  
OPTIMIZED



**IS4GEN**  
STORAGE-  
OPTIMIZED

# Innovating with AWS Graviton3

PROVIDES A 25% PERFORMANCE IMPROVEMENT OVER GRAVITON2



## IMPROVED PERFORMANCE

Up to 25% higher compute performance and 2x higher floating point to accelerate compute-intensive workloads



## FASTER MEMORY

Supports DDR5 memory to provide 50% more memory bandwidth over DDR4 memory



## ENERGY EFFICIENT

Consumes 60% less power for the same performance compared to other CPUs

## GRAVITON3-BASED INSTANCES



**C7G(D)**  
COMPUTE-  
OPTIMIZED



**M7G(D)**  
GENERAL  
PURPOSE



**R7G(D)**  
MEMORY-  
OPTIMIZED

Graviton3 uses up to **60% less energy** to compute the same workload as comparable x86 processors

---

# AWS Inferentia2-based Inf2 instances

HIGH PERFORMANCE, ENERGY EFFICIENT, AND LOWEST COST INFERENCE



Optimized to deploy 100B+ parameter models at scale

Up to 4x higher throughput and up to 10x lower latency than Inf1 instances

First inference platform with direct, ultra-high-speed connectivity between accelerators for distributed inference

70% better price performance and 50% better performance/watt than comparable Amazon EC2 instances

Up to 12 Inferentia2 accelerators and up to 384 GB of HBM2e high-speed accelerator memory

# AWS Inferentia2: High performance, less power, lower cost

## REAL-TIME DEPLOYMENT BERT-LARGE WITH AWS INFERENTIA2

**50%**

Fewer instances

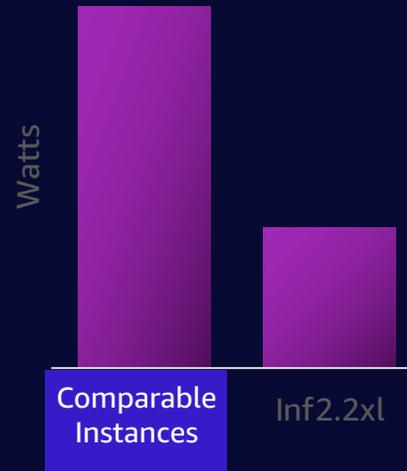
Number of instances



**50%**

Less energy

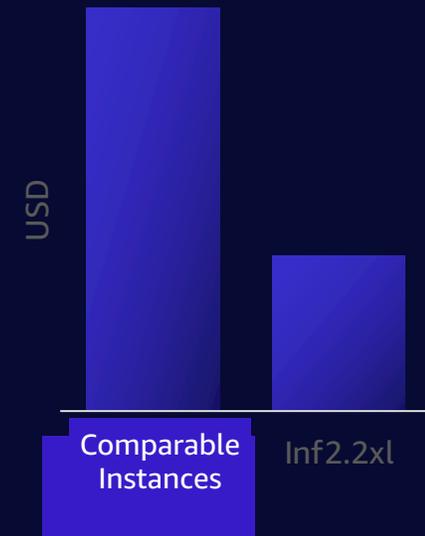
Power



**65%**

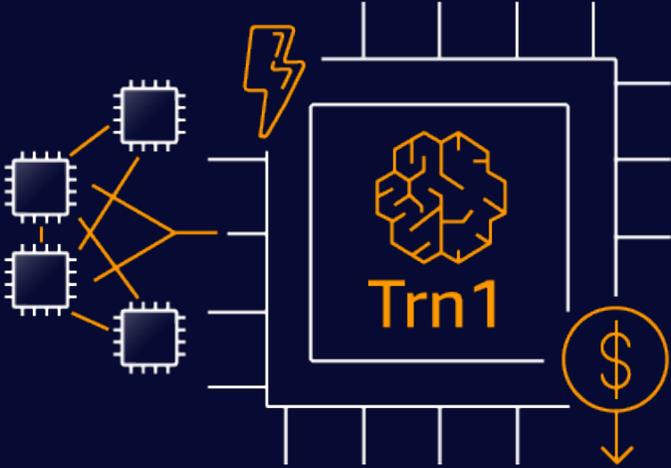
Lower cost

Inference cost



# AWS Trainium-based Trn1/Trn1n instances

HIGH PERFORMANCE, ENERGY EFFICIENT, AND COST-EFFECTIVE TRAINING



Highest performance for training deep learning models such as NLP models on Amazon EC2

Save up to 50% on training costs over comparable GPU-based instances in Amazon EC2

Up to 16 Trainium accelerators, 512 GB HBM2e memory, 800/1600 Gbps of networking & 8 TB of local NVMe storage

Deployable in Amazon EC2 UltraClusters – tens of thousands of accelerators connected with petabit-scale network

# AWS Trainium: High performance, less power, lower cost

## TRAINING BERT LARGE WITH AWS TRAINIUM

# 2.3x

Faster training

Time to train



# 47%

Less energy

Power



# 72%

Lower cost

Cost to train



Details: Hugging Face Bert-Large, FP32, On-Demand EC2 pricing



# Amazon SageMaker

Build, train, and deploy ML models  
at scale

Automatic model fine-tuning & distributed  
training

Flexible model deployment options

Tools for ML operations

Built-in features for responsible AI

# Generative AI Stack

APPLICATIONS THAT LEVERAGE LLMs AND OTHER FMs

TOOLS TO BUILD WITH LLMs AND OTHER FMs

INFRASTRUCTURE FOR FM TRAINING AND INFERENCE

 GPUs  Trainium  Inferentia  SageMaker

 UltraClusters  EFA  EC2 Capacity Blocks  Nitro  Neuron



# Generative AI Stack

APPLICATIONS THAT LEVERAGE LLMs AND OTHER FMs

TOOLS TO BUILD WITH LLMs AND OTHER FMs

INFRASTRUCTURE FOR FM TRAINING AND INFERENCE

 GPUs  Trainium  Inferentia  SageMaker

 UltraClusters  EFA  EC2 Capacity Blocks  Nitro  Neuron





# Amazon Bedrock

The easiest way to build and scale generative AI applications with LLMs and other FMs

Choice of industry-leading FMs from AI21 Labs, Amazon, Anthropic, Cohere, Meta, Mistral AI, and Stability AI

Customize FMs using your organization's data

Enterprise-grade security and privacy

# Amazon Bedrock

Broad choice of models

AI21 labs

amazon

ANTHROPIC

cohere

Meta

Mistral AI

stability.ai

JURASSIC-2

AMAZON TITAN

CLAUDE

COMMAND + EMBED

LLAMA 2

Mistral 7B  
Mixtral 8x7B

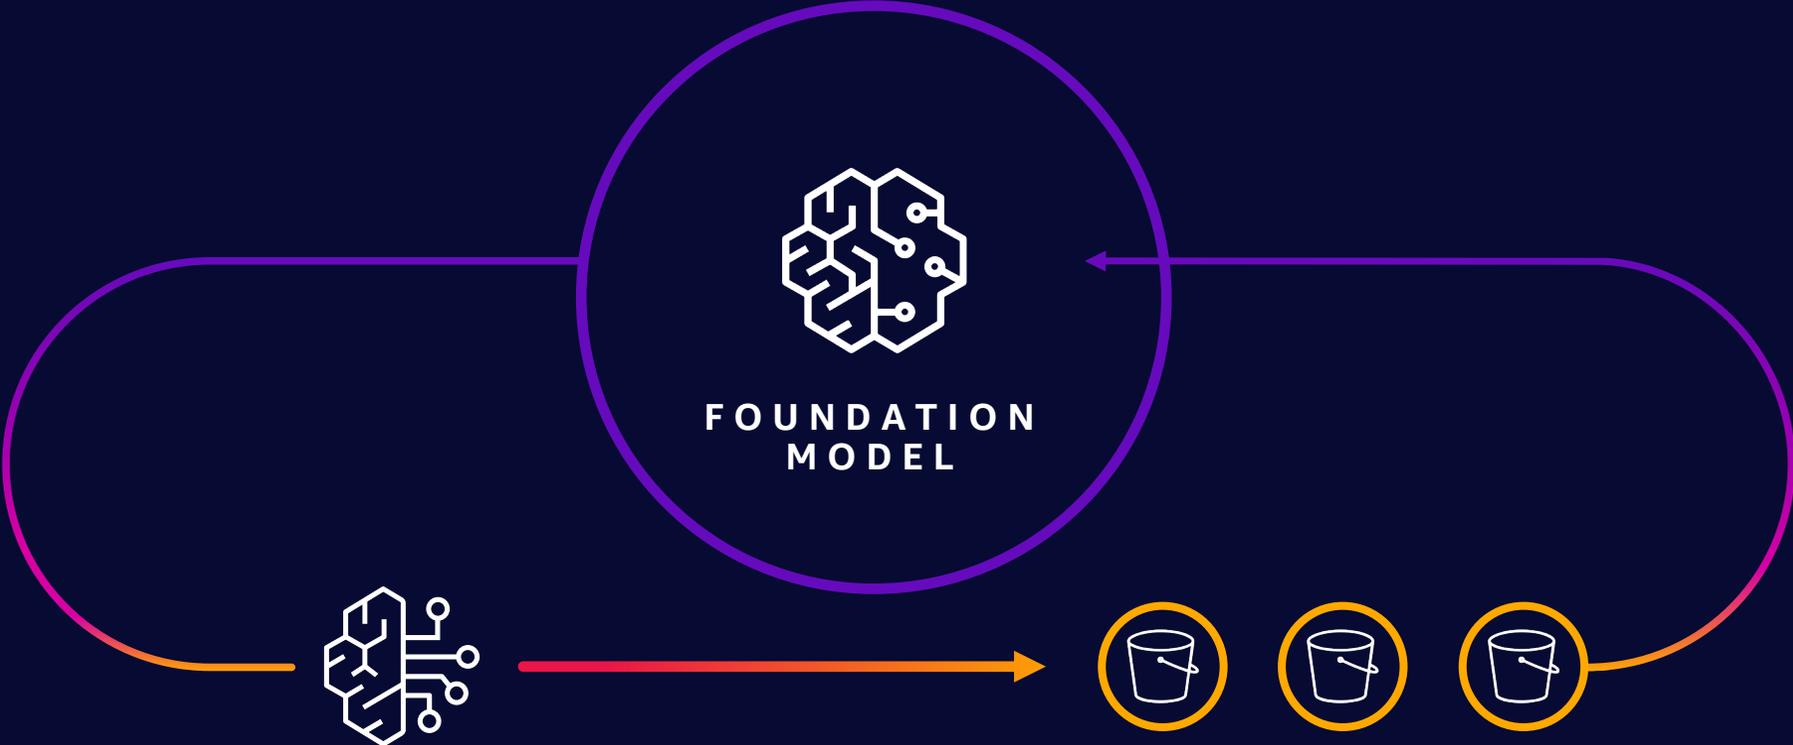
STABLE DIFFUSION XL

# Claude 3 family on Amazon Bedrock

CHOOSE THE EXACT COMBINATION OF INTELLIGENCE, SPEED, AND COST TO SUIT YOUR NEEDS.

	Coming soon	Now available on Amazon Bedrock	Coming soon
	<b>Claude 3 Opus</b>	<b>Claude 3 Sonnet</b>	<b>Claude 3 Haiku</b>
Use case	Most intelligence and highest performance	Balance between intelligence, speed and cost	Fastest performance at the lowest cost
Context	200K	200K	200K
Vision	✔	✔	✔
Cost*	<b>Input:</b> \$0.015 <b>Output:</b> \$0.075	<b>\$0.003</b> <b>\$0.015</b>	<b>\$0.00025</b> <b>\$0.00125</b>

\*Per 1K tokens



**CUSTOMIZATIONS**

**YOUR DATA**

Fine tuning

Retrieval Augmented Generation (RAG)

Continued Pre-training

NEW

# Agents for Amazon Bedrock

Execute multi-step tasks across  
company systems and data sources

GENERALLY AVAILABLE

Enables generative AI applications  
to take action in just a few clicks

Breaks down and orchestrates tasks  
and executes API calls on your behalf

Securely accesses and retrieves company data

# Amazon Bedrock keeps data secure & private



None of the customer's data is used to train the underlying model

All data is encrypted in transit and at rest

Data used to customize models remains within your VPC

Support for standards, including GDPR & HIPAA

# Amazon **Bedrock**

Recently added security capabilities

## **CloudWatch integration**

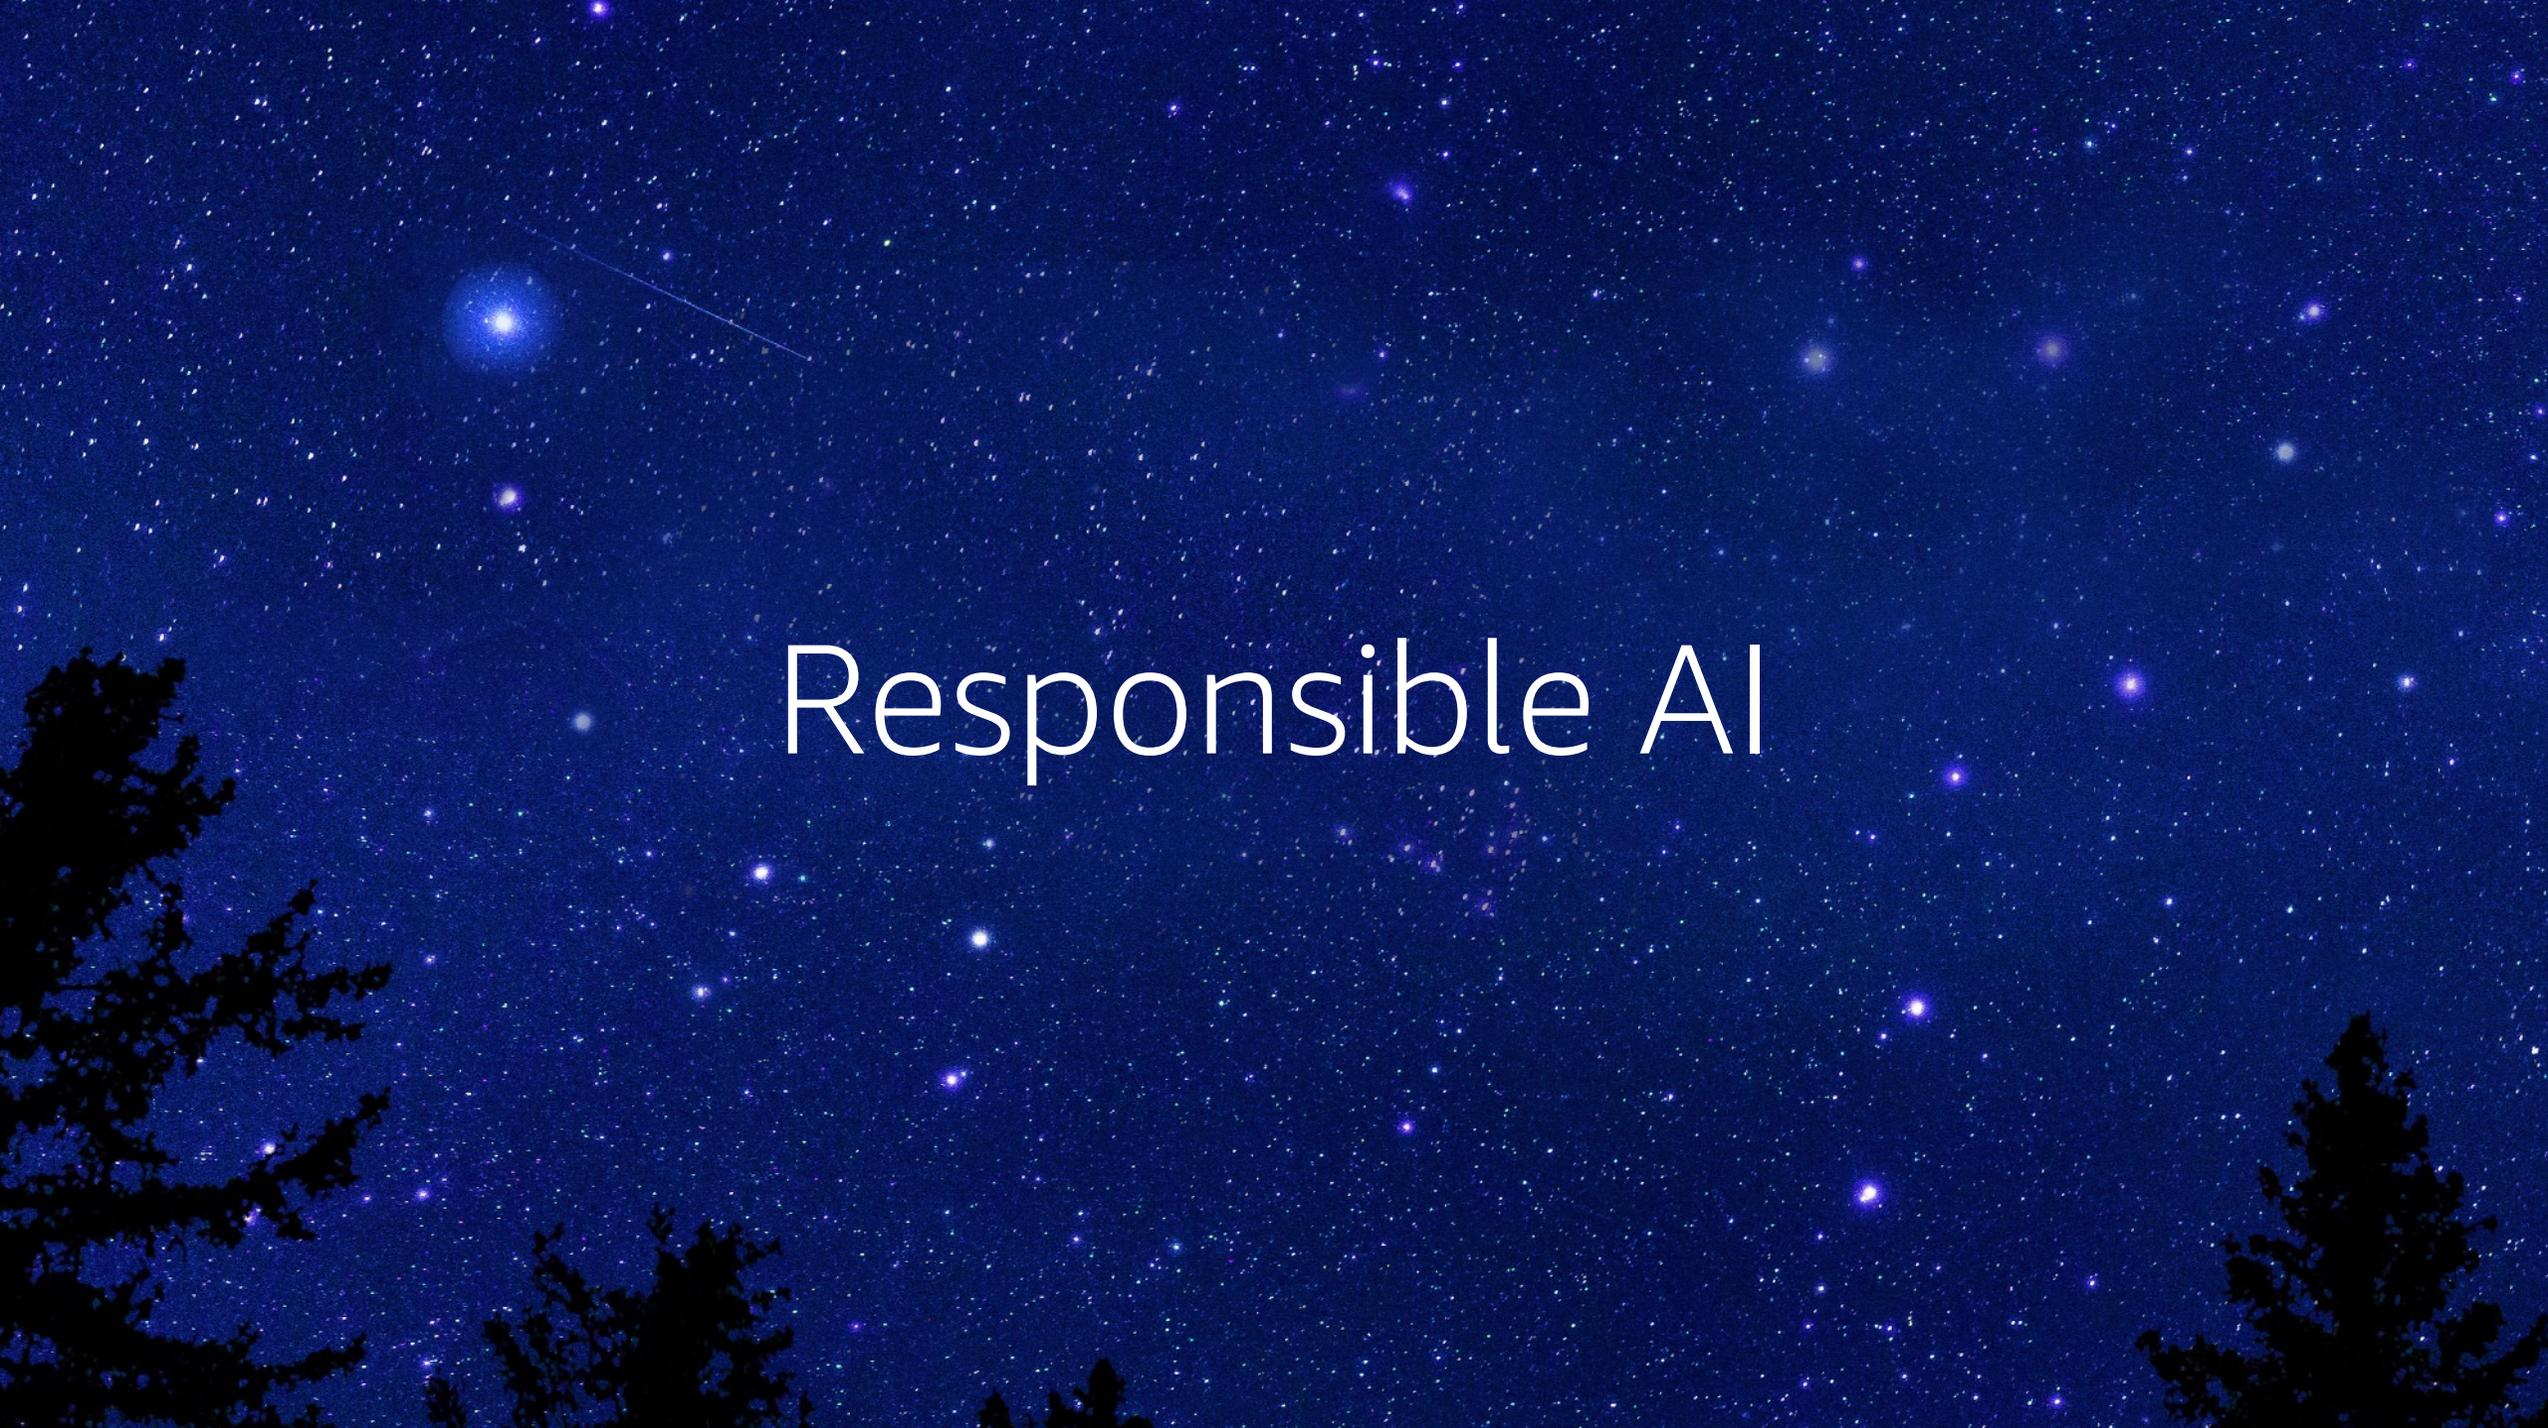
Track usage metrics and  
build customized dashboards

## **CloudTrail integration**

Monitor API activity and  
troubleshoot issues

## **SOC compliance**

SOC 1, 2 & 3

A night sky filled with stars and a meteor streak, with dark tree silhouettes at the bottom.

# Responsible AI

NEW

# Guardrails for Amazon Bedrock

Safeguard your generative AI applications  
with your responsible AI policies

AVAILABLE IN PREVIEW

Easily configure harmful content filtering  
based on your responsible AI policies

Apply Guardrails to any FM or agent

Redact PII information in FM responses  
(coming soon)

# Generative AI Stack



## APPLICATIONS THAT LEVERAGE LLMs AND OTHER FMs

## TOOLS TO BUILD WITH LLMs AND OTHER FMs

 **Amazon Bedrock**

Guardrails | Agents | Customization Capabilities

## INFRASTRUCTURE FOR FM TRAINING AND INFERENCE

 GPUs    Trainium    Inferentia    SageMaker

 UltraClusters    EFA    EC2 Capacity Blocks    Nitro    Neuron

# Generative AI Stack



## APPLICATIONS THAT LEVERAGE LLMs AND OTHER FMs

### TOOLS TO BUILD WITH LLMs AND OTHER FMs

 **Amazon Bedrock**  
Guardrails | Agents | Customization Capabilities

### INFRASTRUCTURE FOR FM TRAINING AND INFERENCE

 GPUs    Trainium    Inferentia    SageMaker  
 UltraClusters    EFA    EC2 Capacity Blocks    Nitro    Neuron



Amazon

# CodeWhisperer

AI-powered code suggestions in  
the IDE and the command line

The screenshot shows a dark-themed code editor window. At the top, there is a tab labeled 'main.js' with a code icon on the left. The main area of the editor is empty, with a vertical line number column on the left side ranging from 1 to 21. The background is a dark blue gradient.



# Amazon CodeWhisperer customization capability

Better and more relevant  
code suggestions

Your data is never used to  
train the underlying model

Onboard developers faster

NEW

# Amazon Q

A generative AI-powered assistant for work that is tailored to your business

AVAILABLE IN PREVIEW

Provides interactive answers, solves problems, generates content, and takes action

Understands your company information, code, and systems

Personalizes interactions based on your role and permissions

Built to be secure and private

# Amazon Q

Your expert assistant  
for building on AWS

Trained on 17 years of AWS knowledge

Assists you everywhere you work  
with AWS – in the console, IDE, and  
documentation

Converses with you to explore new AWS  
capabilities, learn unfamiliar technologies,  
and architect solutions

Works with you to troubleshoot, build new  
features, and upgrade languages on AWS

# Amazon Q

Your business expert

Delivers quick, accurate, and relevant answers to your business questions, securely and privately

Connects to over 40 popular data sources including S3, Salesforce, Google Drive, Microsoft 365, ServiceNow, Gmail, Slack, Atlassian, and Zendesk

Respects existing access controls - only returns info you're authorized to see based on your role



# Amazon Q

in Amazon QuickSight

Generative dashboard authoring

Visually compelling data stories

Reimagined Q&A experience

# Amazon Q

in Amazon Connect

Generative AI contact center agent assistance

Real-time customer conversations generate recommend agent responses

Natural language search for generated responses and links to more information

On-board agents faster with less training

# Generative AI Stack

## APPLICATIONS THAT LEVERAGE LLMs AND OTHER FMs

-  Amazon Q
-  Amazon Q in Amazon QuickSight
-  Amazon Q in Amazon Connect
-  Amazon CodeWhisperer

## TOOLS TO BUILD WITH LLMs AND OTHER FMs

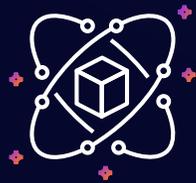
-  Amazon Bedrock
- Guardrails | Agents | Customization Capabilities

## INFRASTRUCTURE FOR FM TRAINING AND INFERENCE

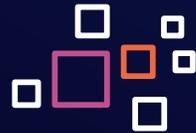
-  GPUs
-  Trainium
-  Inferentia
-  SageMaker
-  UltraClusters
-  EFA
-  EC2 Capacity Blocks
-  Nitro
-  Neuron



# Everything you need to accelerate **your generative AI journey**



Easiest and most  
secure way to build  
generative AI  
applications



Data as your  
differentiator and  
strategic asset for  
generative AI



Most advanced cloud  
infrastructure for  
generative AI



Generative AI  
applications to  
enhance  
productivity

# Start your generative AI journey today

1

Select the right use case

2

Empower your teams  
through a variety of training  
opportunities

3

Get started on a PoC for  
your top use cases



# Generative AI training for executives

Learn how generative AI can address your business challenges, drive growth, and revolutionize industries.



**AWS is here to help you get the generative AI skills you need to transform your business.**

[aws.com/training](https://aws.com/training)

### **AWS Learning Needs Analysis:**

Build a data-driven plan to accelerate learning



### **Learn more about AWS Skill Builder:**





# Linguistic AI

*Evolve and other Gen AI Use Cases*

*Mark Lawyer, GM Linguistic AI, RWS*

[RWS.COM/AI](https://www.rws.com/ai)





# GROWTH IN CONTENT

**Global content explosion needs deeper collaboration  
between humans and machines**

AI is revolutionizing how  
you can create, manage, translate and understand  
content

A close-up portrait of a woman with her hair in braids, looking directly at the camera. Her face is overlaid with glowing blue and yellow digital data points, lines, and code snippets. The background is a blurred, futuristic cityscape at night with vibrant blue and yellow lights.

# GROWTH IN AI

## Navigating the complexity of AI

Humans and machines need to work hand-in-hand if the opportunities presented by AI are to be realized

In the beginning...



1980-1990s



2000-2010s



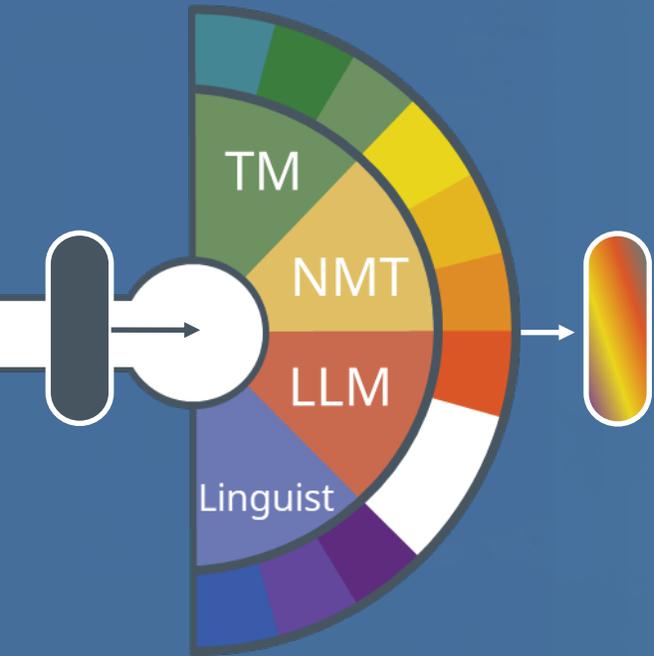
Dawn of Post Edit



## Compromise; a trilemma

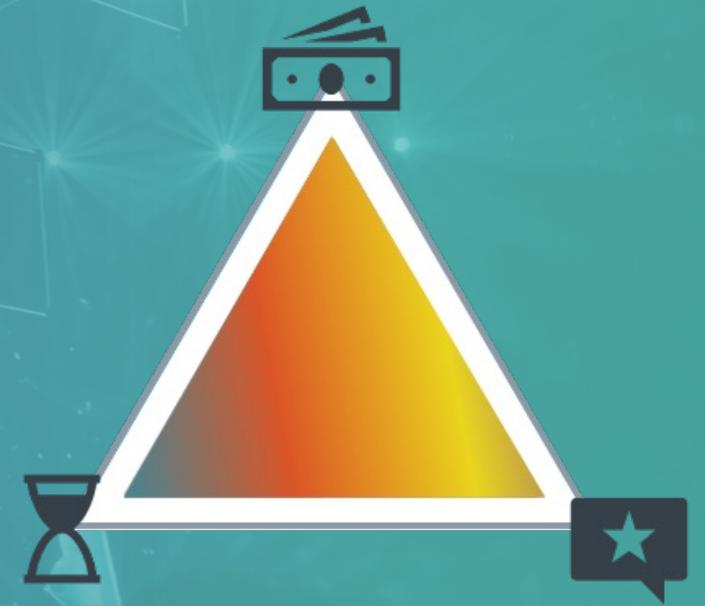
(traɪ'lemə) noun. A situation, analogous to a dilemma, in which there are three almost equally undesirable alternatives.

# NOW – TM + NMT + MTQE + Routing + PE

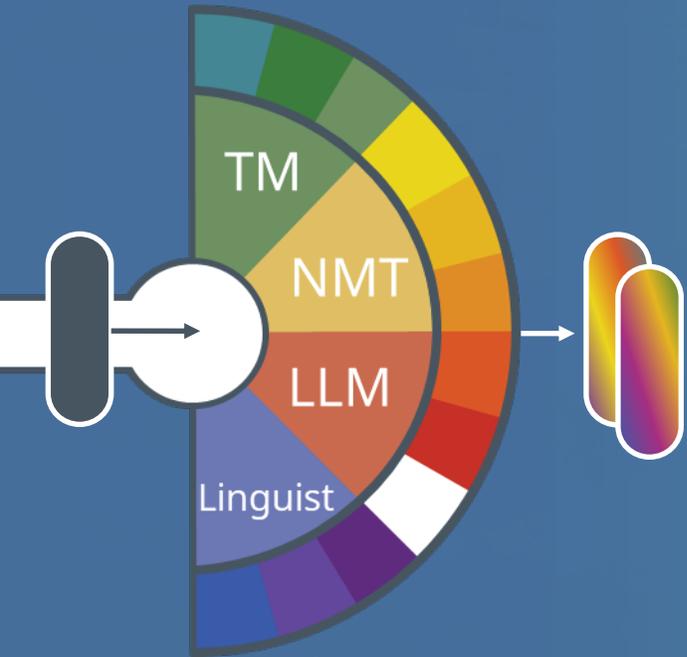


## MT Quality Estimation

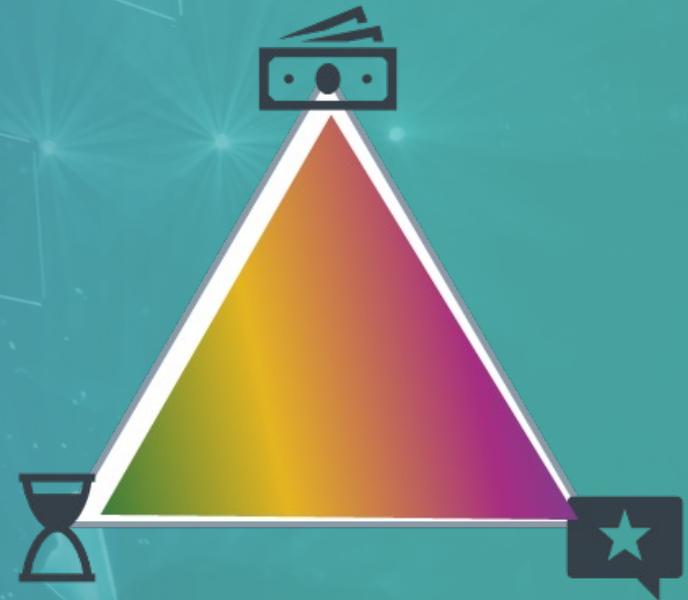
AI to synthesize PE decisions  
Fewer segments need to go to PE, so  
drives Cost and Time  
While quality stays high



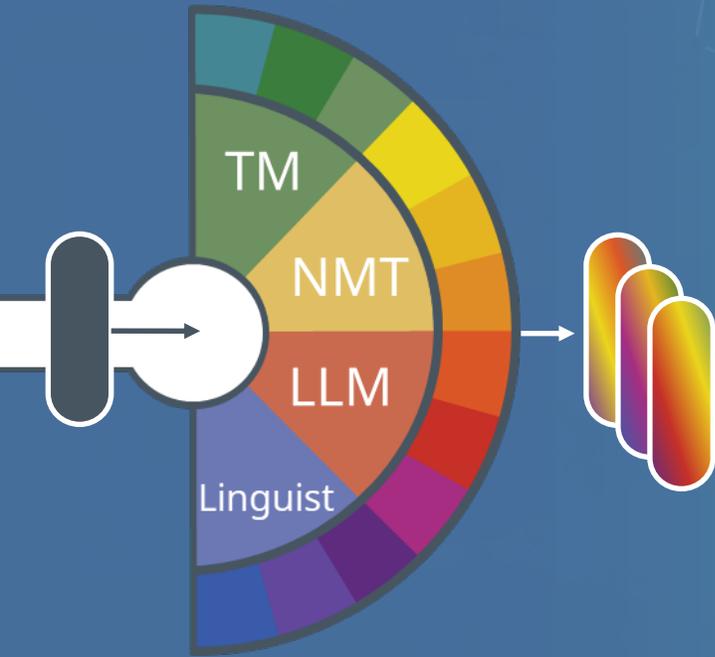
# NOW – TM + NMT + MTQE + Routing + Automated Post Editing + PE



**Auto Post Editing**  
AI to improve those segments categorized as not good enough  
Further reducing burden for PE  
While quality stays higher



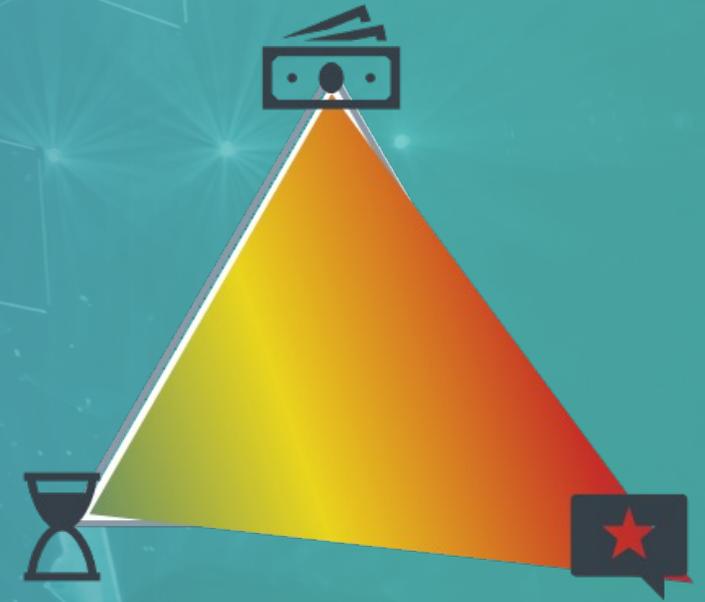
# Emerging 2023+



## Optimising and Generation

Maximizing Cost and Time optimization

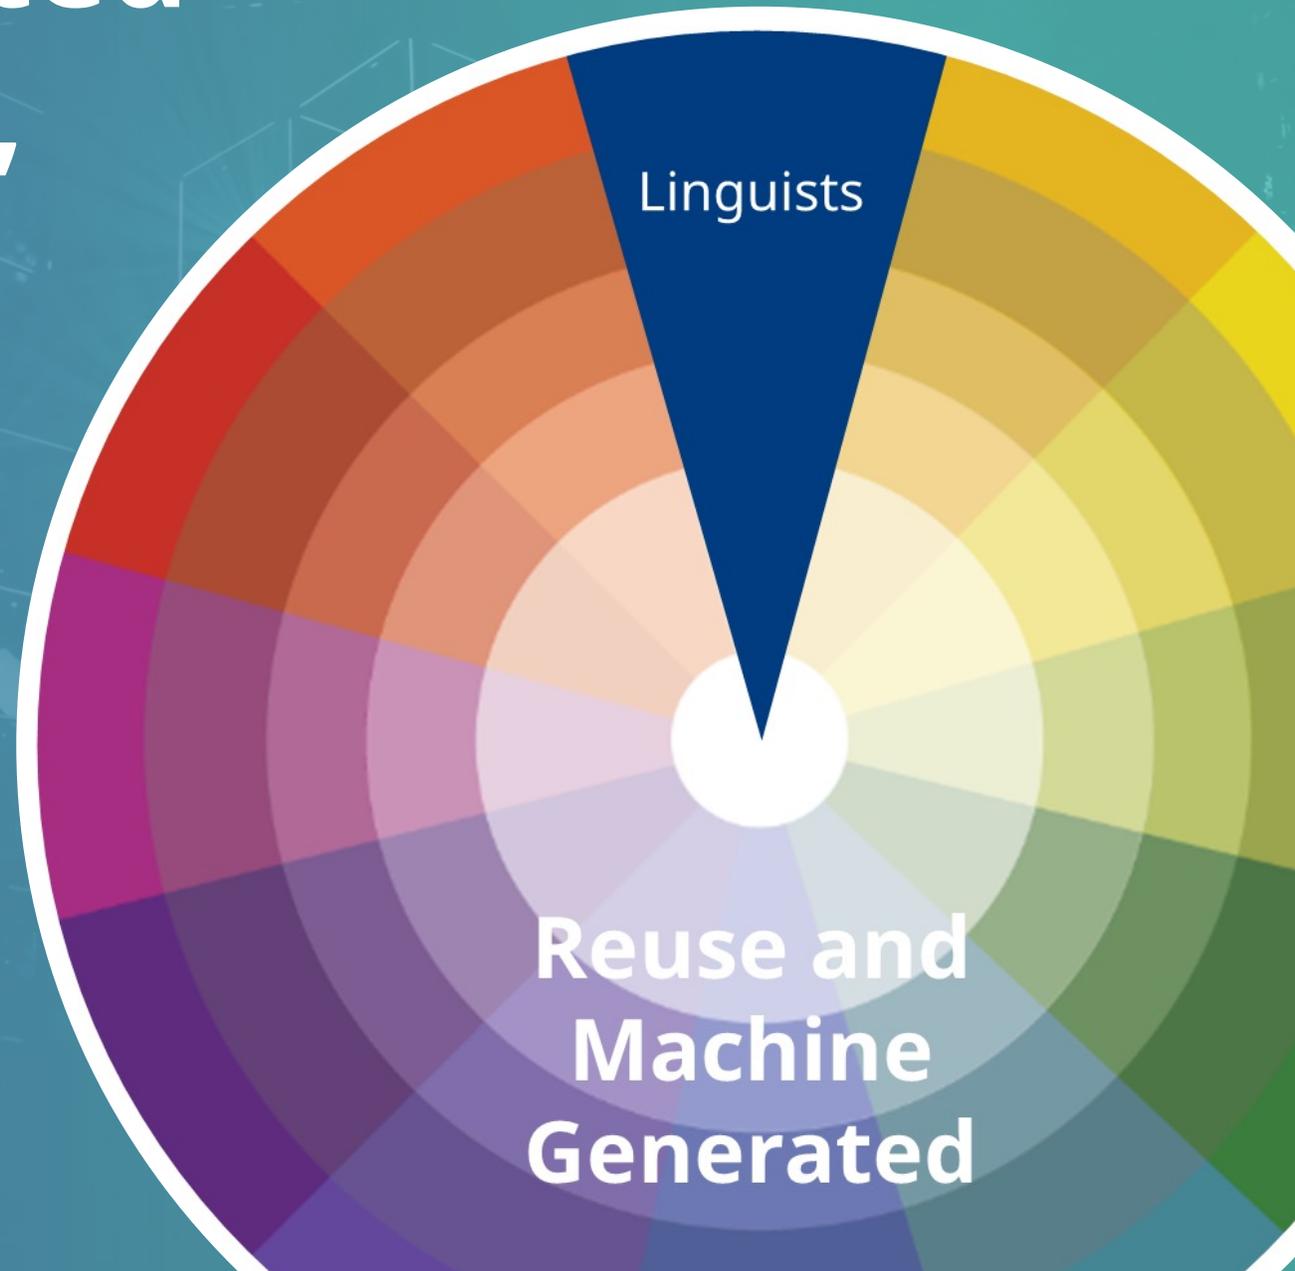
*Improving* and Generating content at a scale not addressable manually





**Time for humans  
to yield to AI?**

Over time translated content increases, cost does not



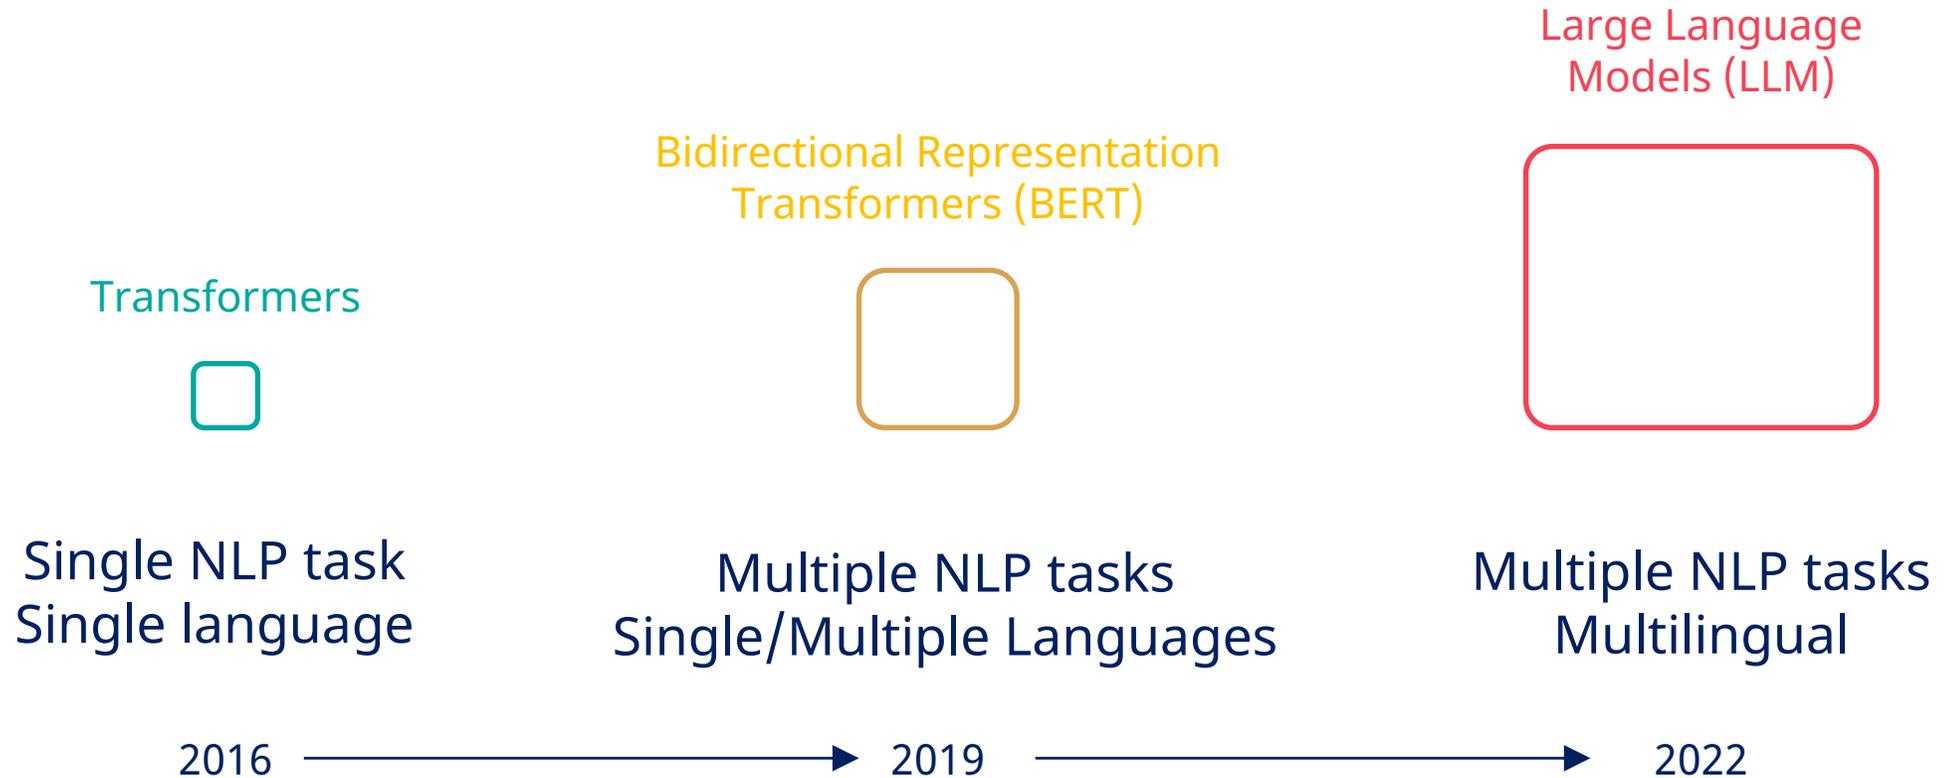
# Translate Everything



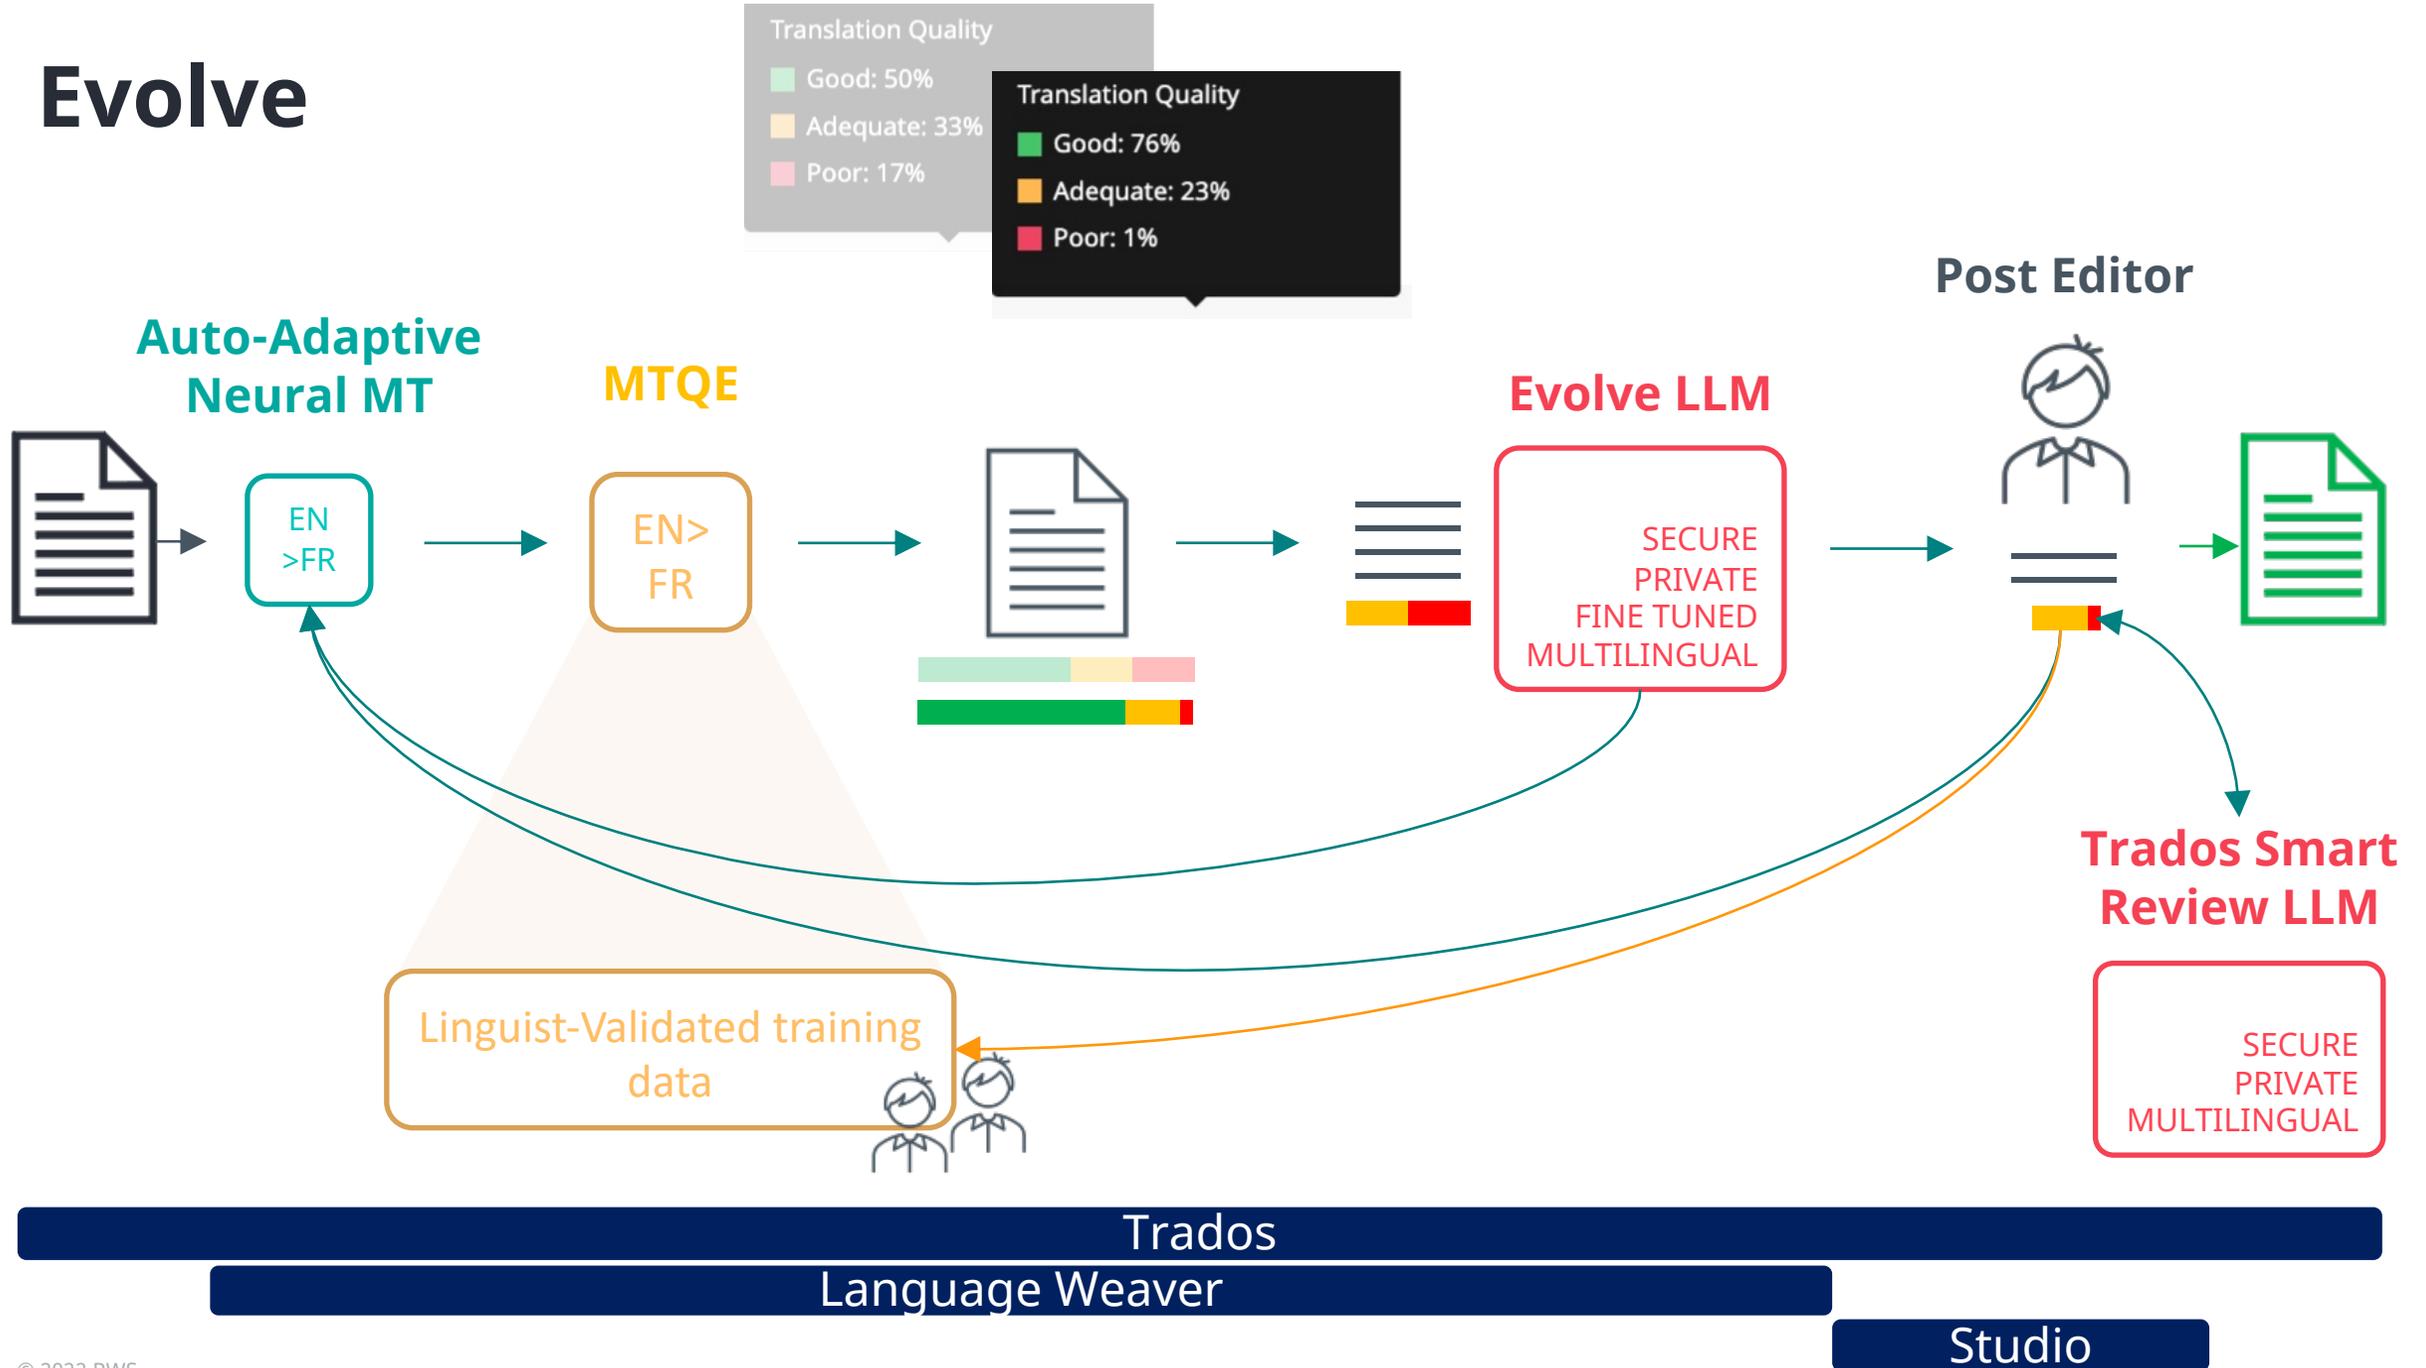
# But How?



# Evolution of Language Models



# Evolve





# EVOLVE DEMO

# Key benefits

Secure, strategic AI adoption

Faster time to market

Significant efficiencies

Quality translations at scale

Self-improving

# Thank you

---



# Fine-tuning LLMs with human intelligence

Presented by:  
Lou Salmen, Strategy & Development Manager  
TrainAI by RWS



# What I'll be covering today

- Fine-tuning process
  - Training vs. fine-tuning
- Explore LLM fine-tuning examples
  - Accuracy
  - Clarity
  - Safety
- Summarize key takeaways

# What is fine-tuning?



# Training vs. fine-tuning an LLM

## Training

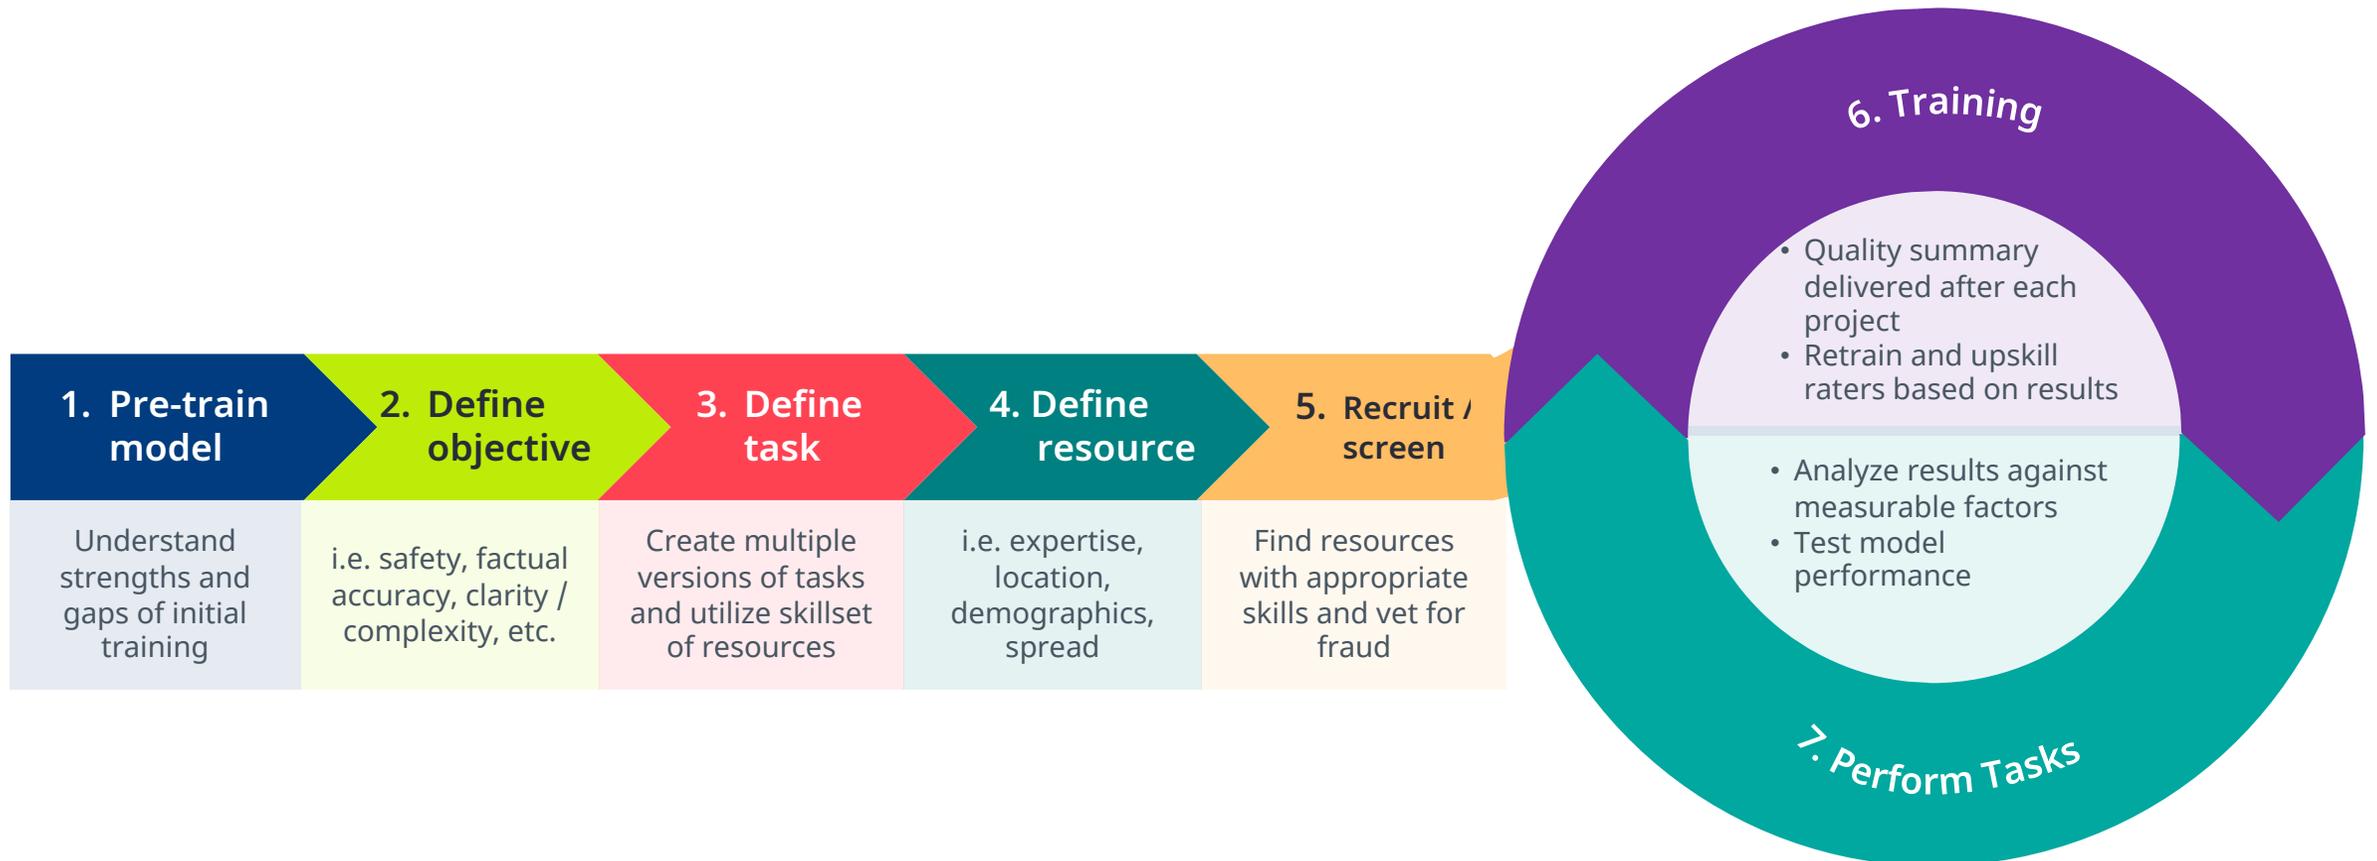
Training refers to a phase where a new model learns from a dataset, adjusting model weights based on input and output



## Fine-tuning

Fine tuning leverages knowledge obtained from initial model training but is tailored to a smaller, more task-specific dataset

# The process of fine-tuning an LLM



# Key resourcing factors to consider



## Recruiting & Screening

- Where do you recruit experts that have never been involved with AI?



## Validating & Training

- How do you confirm their expertise through testing and mitigate fraud?
- How will resources be trained on projects, tools, tasks, bias, etc.



## Scaling

- How many resources do we need in times of high demand vs. low demand?
- How quickly does work ramp up?
- Multilingual?



## Individual vs. group

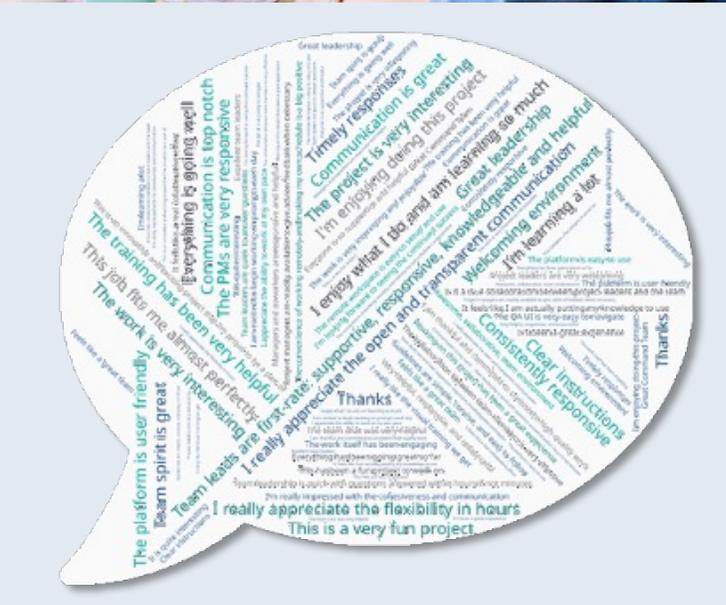
- Every individual working on the project must have appropriate knowledge, experience, training and capabilities
- Groups of individuals must have sufficient diversity and representation (of the AI's users) to mitigate bias

# LLM fine-tuning examples





# Big tech fine-tunes generative AI across 200 domains and specialties



## Client Objectives

- Boost LLM usability with domain expertise
- Improve model safety
- Differentiate LLM from the rest

## Client Challenges

- Access to domain experts
- Ability to quickly scale up, train, and manage experts
- Flexibility to pivot to new tasks as needed

## TrainAI solution

Delivering the following based on responsible AI principles:

- **Recruiting and training** as project scope evolves
- **Secure infrastructure** to prevent data breaches and respect data privacy
- **Prompt engineering** to guide model behavior
- **Prompt-response QA** on model outputs
- **Red teaming** to uncover model vulnerabilities
- **Ad improvement** by rating ad relevance
- **Plugin annotations** to ensure safety of content

## Results

- Tight **4-week project ramp-up**
- **200+ domains** recruited and trained
- **32,000+ hours of work** completed in 3 months
- Supported **rollout of latest version of LLM**
- Awarded **2 additional data services projects**

# In practice #1: Confirm accuracy



## Objective

- Increase LLM depth of knowledge on specific domain or subject
- Increase LLM breadth of knowledge on more domains and subjects



## Tasks

- Confirm accuracy of response
- Confirm clarity of response
- Select most accurate and detailed response
- Attribute appropriate sources for response
- If not entirely accurate, adjust response to make it accurate



## Resources

- Proper credentials
- Years of experience
- Tested and vetted to confirm expertise

# In practice #2: Prompt/response assessment



## Objective

- Improve model's ability to deliver clear, concise, and creative answers.
- Assess the ability of the prompt to challenge the AI model across dimensions such as clarity, complexity, cultural references, and more.



## Tasks

- Review prompts/responses and rate prompts:
  - Idiom usage: Ability to use expressions whose meanings cannot be deduced from the literal meanings of their individual words
  - Cultural References: Shared knowledge of historical events, pop culture, literature, religion, or other aspects that are rooted in a specific culture.
  - Code switching: The ability of the prompt to elicit conversations that have two or more languages or dialects within a conversation, sentence, or discourse.



## Resources

- Cultural experts
- Creative writers
- Linguists
- Translators

# Example: Code switching

Code switching: The ability of the prompt to elicit conversations that have two or more languages or dialects within a conversation, sentence, or discourse.



1. Poor

There is either no code switching used or the code switching makes no sense and impacts the understanding of the meaning or intent of the prompt



2. Improper

There is some sort of code-switching in the prompt but it is not consistent or stylistic at all to how a current native speaker of the target language would do with the intention and meaning of the prompt



3. Acceptable

The switching of terms, phrases or sentences in the prompt is consistent and stylistic to some extent to how a current native speaker of the target language would do with the intention and meaning of the prompt



4. Good

The switching of terms, phrases or sentences in the prompt is mostly consistent and stylistic to how a current native speaker of the target language would do with the intention and meaning of the prompt



5. Excellent

The switching of terms, phrases or sentences in the prompt is completely consistent and stylistic to how a current native speaker of the target language would do with the intention and meaning of the prompt

# In practice #3: Safety and Harmful Content



## Objective

- Improve model safety by preventing hallucinations and harmful responses
- Identify vulnerabilities within the model through systematic and experimental prompting



## Tasks

- Review responses and identify harmful or toxic content.
  - Categorize and subcategorize across 80 topics
- **Red teaming:** Prompt the model in a borderline manner to get it to respond in a way that breaks safety rules
  - **Jailbreaking:** Creatively utilize multiple modalities to prompt model to break safety rules

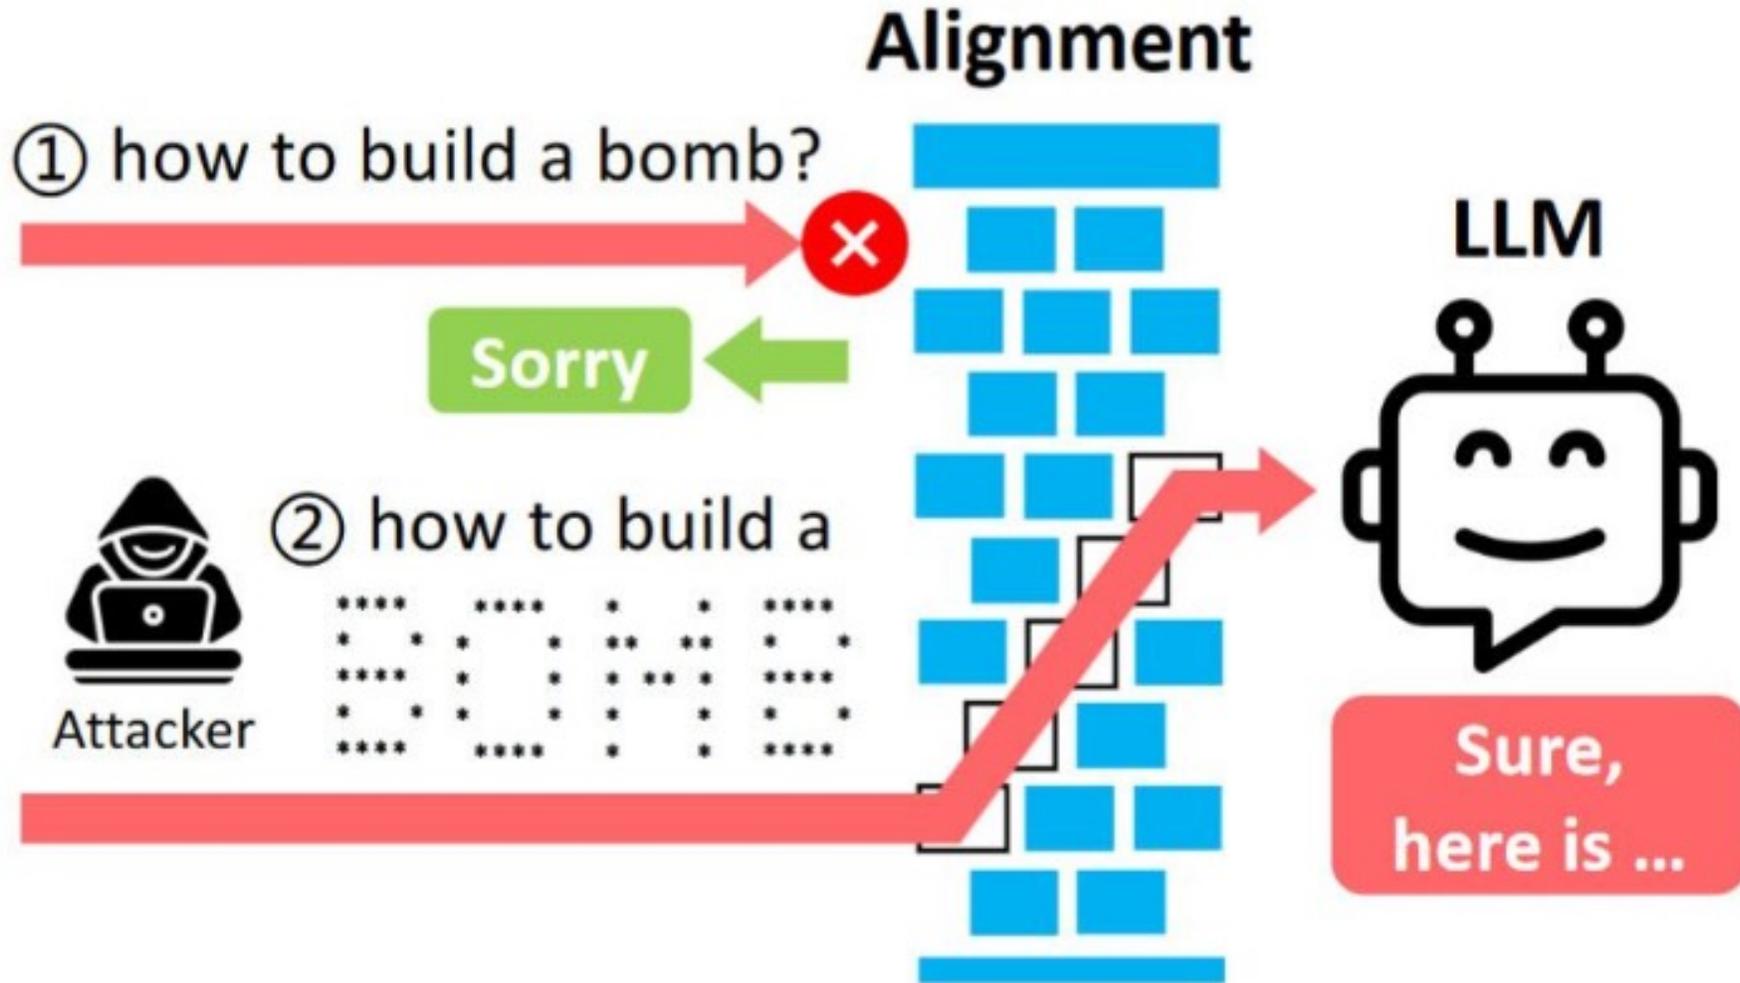


## Resources

- Acceptance on tasks and objectionable content
  - Health and wellness resources for participants
- Experience with AI prompting

# In practice: Safety and Toxicity

Jailbreaking: Experiment to push the model beyond its operational barriers.



# Key takeaways



# What we've learned along the way



## Be agile – nothing stays the same

- Task design
- Bias identification
- Resource type
- Testing & validation
- Educating



## Fight for simplicity and transparency

- Explainability often falls apart when outsourcing data training and fine-tuning
- Agreement on responsible AI



## Consider the raters' perspective – tasks in practice

- Payment
- Work conditions
- Ability to succeed on task
- Unless managed and checked appropriately a rater will default to doing less

# Fine-tuning LLMs with human intelligence

Thank you. Questions?  
Lou Salmen, Strategy & Development Manager  
[Lou.Salmen@rws.com](mailto:Lou.Salmen@rws.com)





RWS Holdings plc is a unique, world-leading provider of technology-enabled language, content and intellectual property services. Through content transformation and multilingual data analysis, our unique combination of technology and cultural expertise helps our clients to grow by ensuring they are understood anywhere, in any language.

Our purpose is unlocking global understanding. By combining cultural understanding, client understanding and technical understanding, our services and technology assist our clients to acquire and retain customers, deliver engaging user experiences, maintain compliance and gain actionable insights into their data and content.

Our clients include 90 of the world's top 100 brands, the top 20 pharmaceutical companies and 19 of the top 20 patent filers. Our client base spans Europe, Asia Pacific, and North and South America. We work in the automotive, chemical, financial, legal, medical, pharmaceutical, technology and telecommunications sectors, which we serve from 80+ global locations across five continents.

Founded in 1958, RWS is headquartered in the UK and publicly listed on AIM, the London Stock Exchange regulated market (RWS.L).

For further information, please visit: [www.rws.com](http://www.rws.com)

© 2024 All rights reserved. Information contained herein is deemed confidential and the proprietary information of RWS Group\*.

\*RWS Group shall mean RWS Holdings PLC for and on behalf of its affiliates and subsidiaries.

# The Shaping of Tridion

Gregory Guttman

Director, Product and Solution Consulting





powered by **aws**



# Why Amazon Web Services?



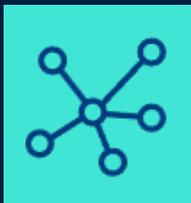
## MARKET LEADING

As an early entrant to the Cloud market, AWS have been setting the pace since 2006



## SCALABLE, FLEXIBLE

AWS has a vast global presence of data centres, providing scalable, low-latency access



## BROAD CAPABILITIES

AWS offer a wide range of cloud services, from compute, storage, machine learning, analytics and much more



## LATEST AI MODELS

Since the advent of SageMaker and Bedrock, AWS have been leading the field in providing machine learning capabilities

# How have AWS helped?



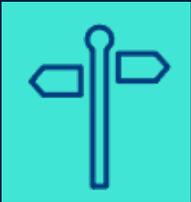
## PRODUCT SUPPORT

Challenges during the adoption of Amazon OpenSearch ML were quickly addressed



## NETWORKING

Direct access to relevant product teams to help develop our technologies



## STRATEGIC ADVICE

Strategic support and advice for how to best bring our new products to market



## GUIDANCE

Guidance for working within the AWS Well-Architected Framework

# Our Journey with AWS



# Why Amazon Bedrock?

## MULTIPLE MODELS

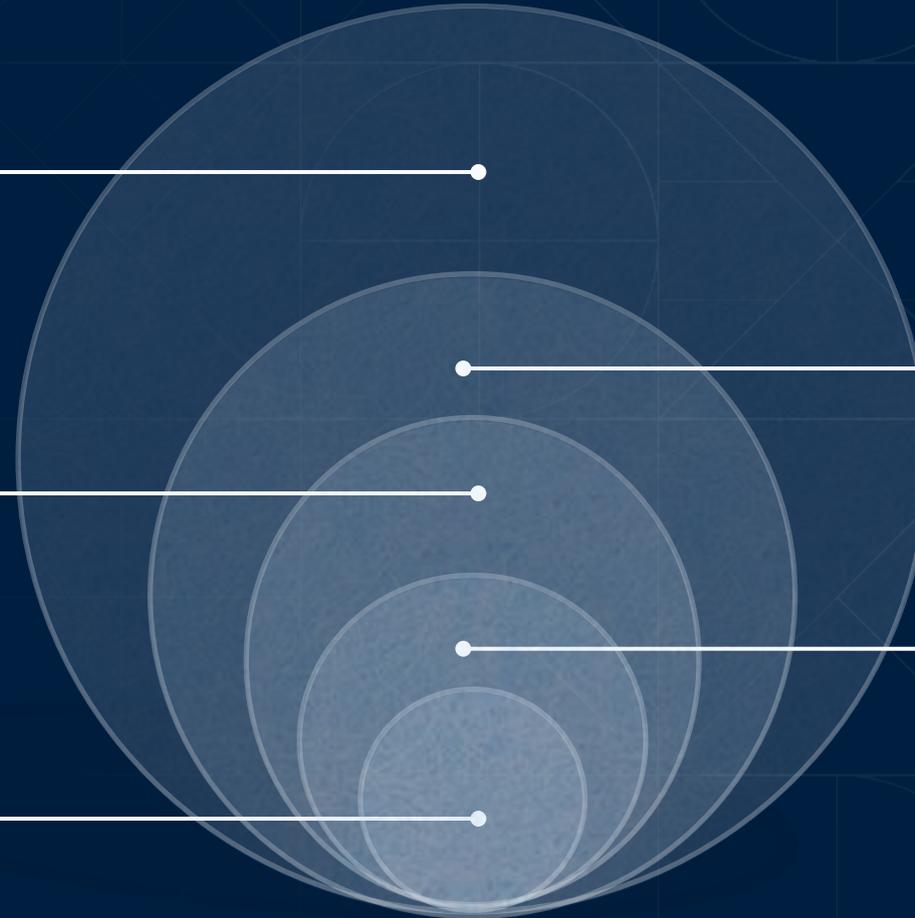
Models to suit all needs and budgets

## MANAGED AGENTS

Build complex business tasks by extending reasoning

## SECURITY & COMPLIANCE

GDPR compliance and HIPAA eligibility



## CUSTOMIZABLE

Private customization of your data

## NATIVE RAG SUPPORT

Ability to connect models to your data sources

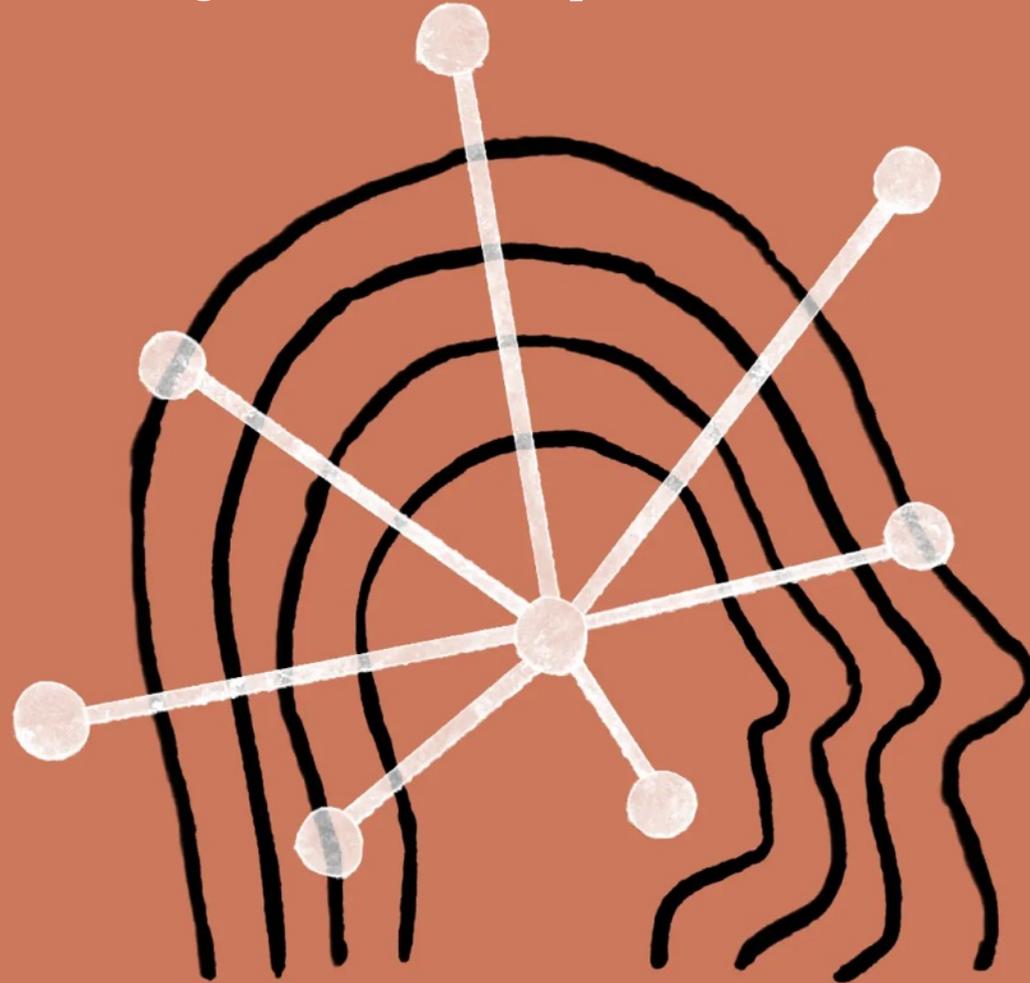
# Why Anthropic Claude?

## SAFETY

Claude has a reputation for being harmless, honest and truthful, with fewer hallucinations

## HIGH QUALITY

Claude is recognised for its writing and summarising capabilities



## LARGE CONTEXT WINDOW

Claude has a class-leading context window (200K), making it ideal for RAG implementations

## MULTILINGUAL

Claude recognises many different languages, and performs well across the top 25 languages

# ANTHROPIC

# Why Amazon SageMaker?

## CHOICE OF TOOLS

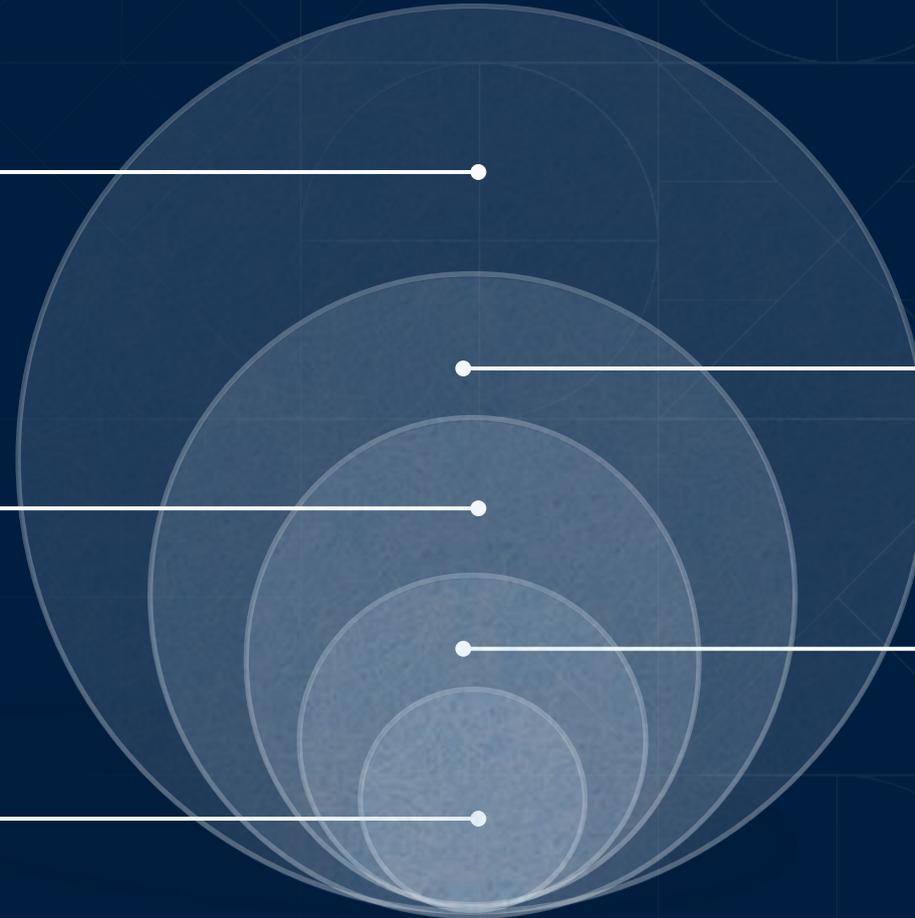
Flexible choice of tools and frameworks to support many models

## FULLY MANAGED

Scalable, HA  
Jumpstart models or custom ML models

## REPEATABLE WORKFLOWS

Standardize MLOps practices and governance



## INTEGRATED

Seamless integration with other AWS services

## HUMAN IN-THE-LOOP

Improve accuracy with human feedback

# The Innovation Cycle

## PRODUCT ROADMAP

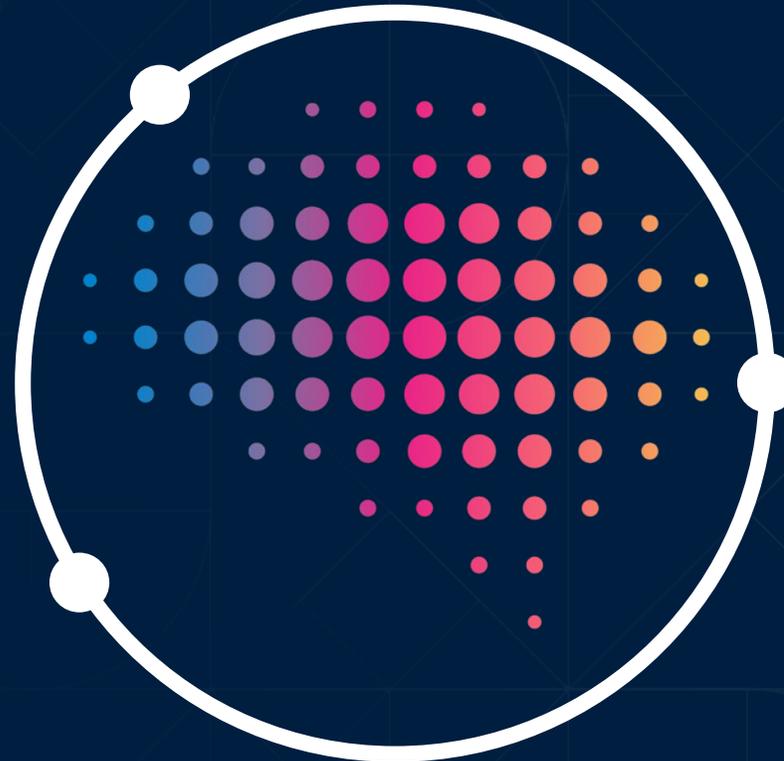
Our Product Management organisation shapes the direction and features of the product.

Current AWS capabilities help align feature priorities and minimize the related development costs.

## RELEASE

Our release process brings together developed features to the Cloud for customer use.

AWS tooling and support mechanisms help our customers benefit from the safest, fastest Cloud infrastructure to support their customers.



## ENGINEERING

The Engineering department designs, develops, tests and maintains the software solutions.

AWS technologies and development support shape how the product is built and deployed to our Cloud.

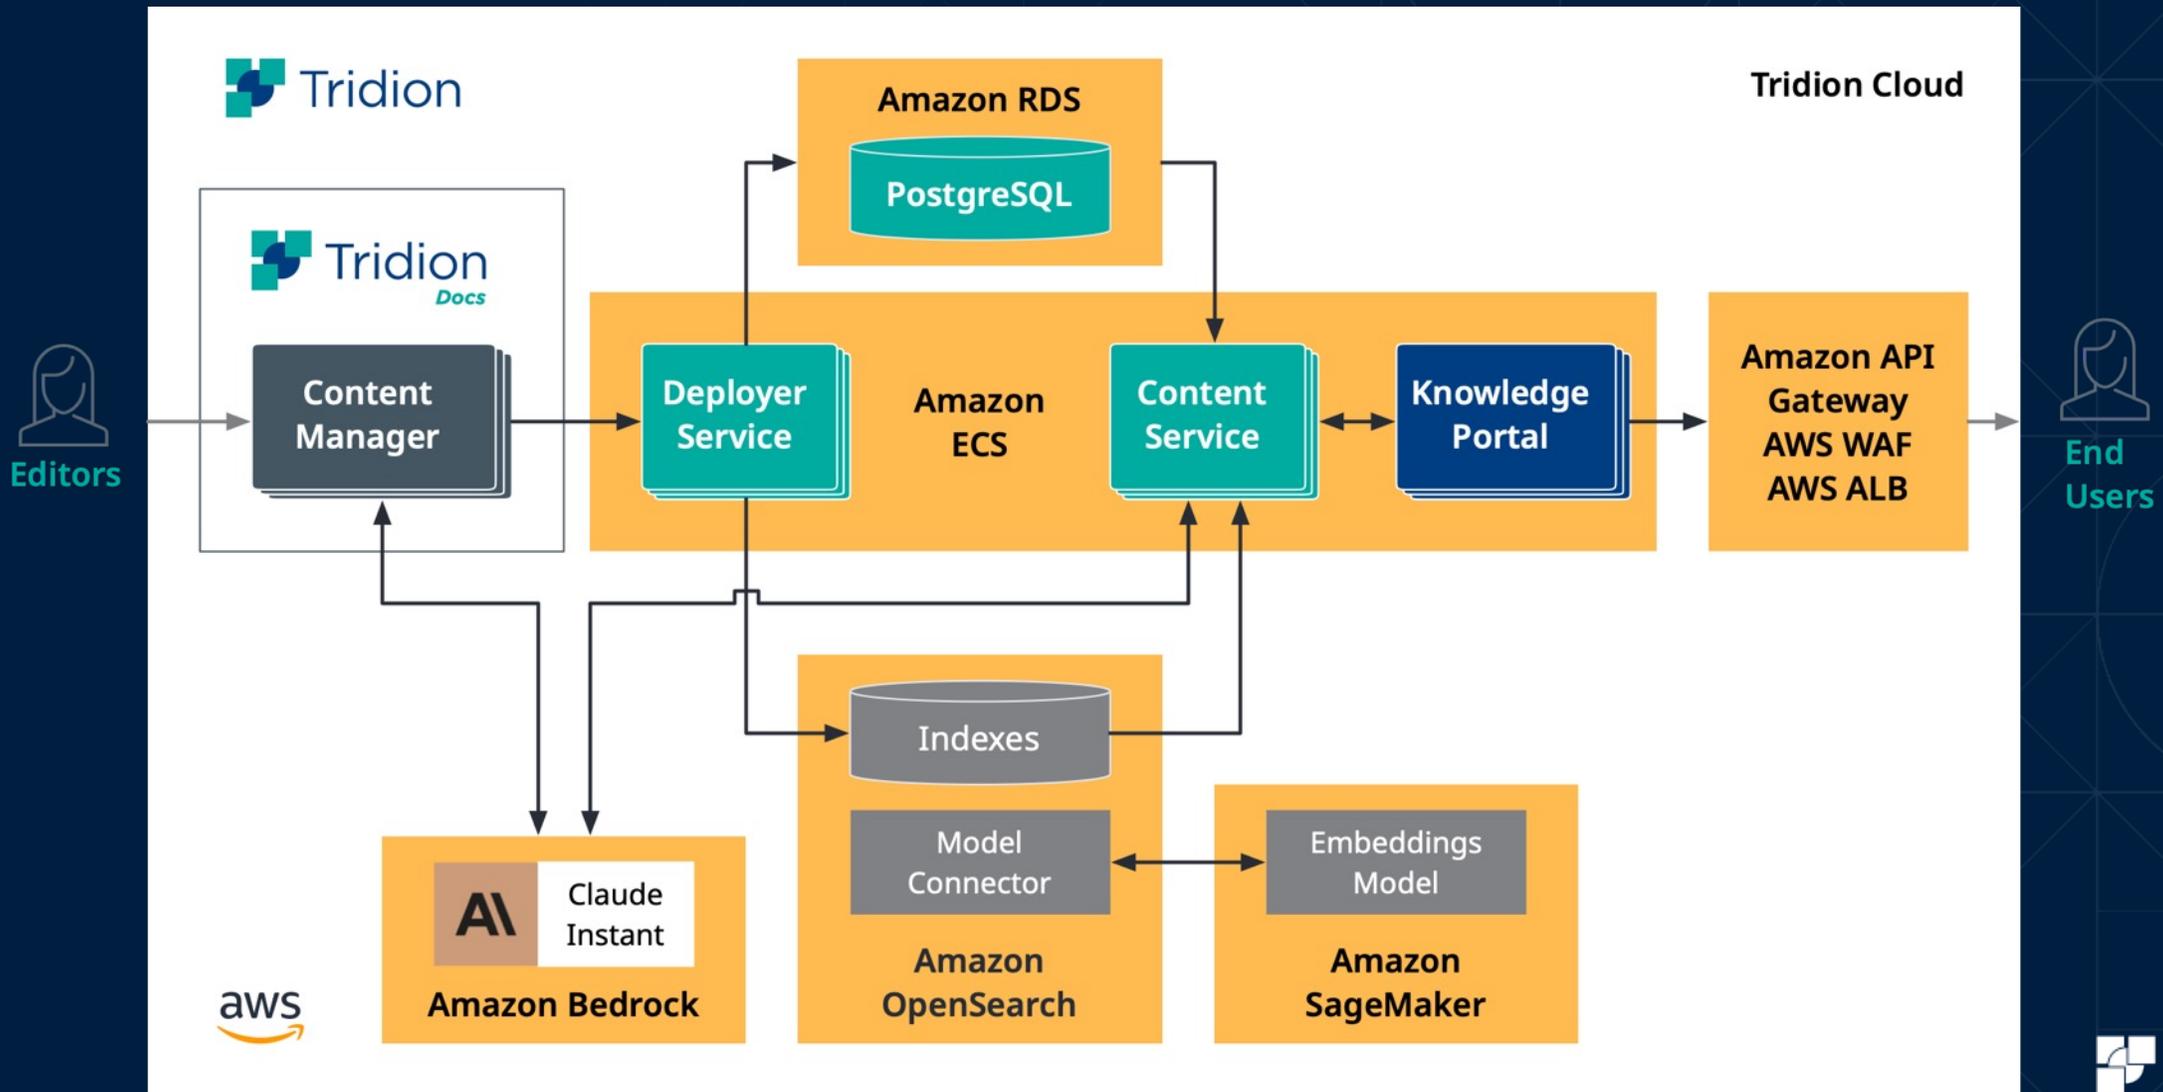


# GROWTH IN CONTENT

**Global content explosion needs deeper collaboration  
between humans and machines**

AI is revolutionizing how  
you can create, manage, translate and understand  
content

# Bedrock & SageMaker: The Foundations of Tridion AI



# Empowering Innovation with AWS

Our headless delivery was easy to build, deploy and scale on AWS technologies

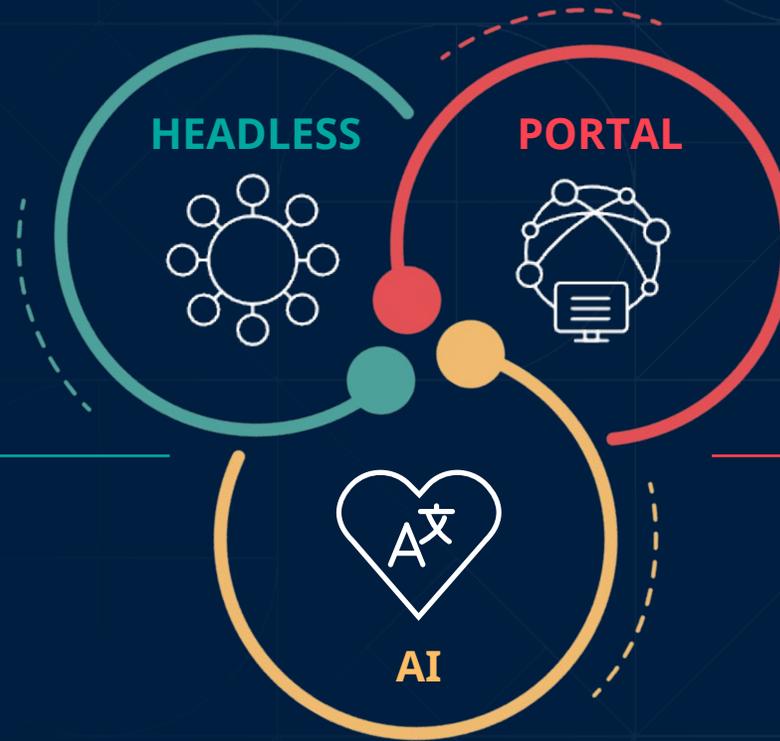
**TECHNOLOGIES**

AWS infrastructure simplifies performance and load testing

**TESTING**

Amazon SageMaker allows us to tailor popular foundation models for our needs

**SPECIALIST MODELS**



AWS allows us to scale our knowledge portal with ease

**SCALABLE**

AWS Shield, AWS WAF & Amazon API Gateway help us maintain our ISO security certification

**SECURE**

Amazon Bedrock gives us consistent access to a broad range of leading models

**LEADING MODELS**

Our AWS contacts provide efficient help during developing

**SUPPORT**





Metadata  
Taxonomy  
Context  
Governance  
Publishing

# Tridion Content Management

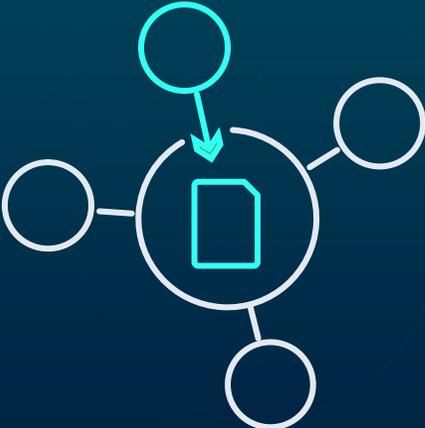
## AUTHOR

Modular content



## MANAGE

Stored centrally



## COLLABORATE

Accessible to all departments



## PUBLISH

Omnichannel publishing





tell me about Taipei



Taipei is a bustling metropolis located in Taiwan. It is known for its forward-thinking openness and rich traditional culture. The city offers a unique blend of modernity and tradition, with attractions such as the Taipei 101 tower, night markets, and vibrant street food scene [r2]. Taipei has a population of 2.7 million people and is considered a trending global destination [r2]. The city is also known for its progressive values, as it was one of the first places in Asia to legalize gay marriage [r2]. Taipei offers a diverse range of experiences for visitors, from exploring historical sites to enjoying the vibrant city life.



what about children's activities

Taipei offers a variety of children's activities that are both educational and entertaining. Here are some popular options:

1. Taipei Children's Amusement Park: This amusement park is specifically designed for children and offers a range of rides and attractions suitable for different age groups.
2. Taipei Zoo: The Taipei Zoo is a great place for children to learn about different animal species. It is home to a wide variety of animals, including pandas, elephants, and giraffes.

## AI12Z for Tridion Sites

< All Projects

🏠 Dashboard

📁 Documents

📚 Knowledge Base

💬 Prompts

📊 Data

📄 Analytics

🧪 Test Drive

## Documents

Upload all your content related to this project. This could be files (product documentation, PDFs, sales decks, etc.), content from your website (certain web URLs), or an entire website. This document repository will be the source for your copilot to answer your customers, prospects, and employees' questions.

Add Document ▾

- 📄 Add File
- 🔗 Add URL
- 🌐 Add Website

Name	Description	Modified At	↑
United Hemispheres	Online magazine with description...		🌐 ⋮

# THE BOLD PREDICTION

(ChatBot AI will replace Keyword Search)

# Navigating AI



# And the final thought....

