



REINVENTING ENERGY

Generative artificial intelligence

Table of contents

AWS for Energy: Helping solve today's energy challenges	3
Energy Systems of the Future: The tools for navigating the road ahead	4
Powering Efficiency and Smarter Thinking	5
Changing the Game With Generative AI	6
Generative AI in Action	8
Optimizing safety and sustainability / Reducing cost and downtime	8
Simulating key energy scenarios / More efficient data management	9
The AWS AI Difference	10
Amazon Bedrock	11
Amazon CodeWhisperer	12
Amazon SageMaker Jumpstart	13
Using AI responsibly	14
The AWS Generative AI Innovation Center	15

ABOUT AWS FOR ENERGY

AWS for Energy

Reinventing the energy industry.

Amazon Web Services (AWS) is the world's most comprehensive and broadly adopted cloud, offering over 200 fully featured services from data centers globally. Millions of customers – including the fastest-growing startups, largest enterprises, and leading government agencies – are using AWS to lower costs, become more agile, and innovate faster.

With today's energy industry rapidly changing, the need for energy companies to innovate is critical. As customer behaviors evolve, energy demand increases, and the need to decarbonize accelerates, energy organizations will need to increase innovation through the latest cloud technologies such as artificial intelligence (AI) and machine learning (ML), edge computing, and many others.

This ebook, focused on generative artificial intelligence (AI), is part of an ongoing series that will dive deep into the critical cloud technologies that are helping energy companies across the value chain reinvent the industry as they build the energy systems of the future.



These energy systems of the future, built with AWS' deepest and broadest cloud services, will result in energy systems that are more:



Resilient

Handles critical events, including increasingly volatile weather, new consumption patterns and consumer expectations, regulatory revisions, supply chain disruptions, and economic volatility.



Responsive

Senses and adapts in real-time to deal with unexpected changes in the operating environment or respond quickly to emergencies.



Resourceful

Maximizes existing assets, people and information, which leads to improved operational excellence and financial performance.



Reliable

Scales across a large footprint sustainably, while ultimately delivering safe, reliable energy to everyone who depends on it.



Renewable

Meets today's demand while making the shift to renewable energy sources and contributing to a balanced global energy mix.

INTRODUCTION

The Energy Systems of the Future

The tools for navigating the road ahead.


Working together with an expansive network of energy partners, customers and startups, AWS is building the energy systems of the future and helping leading organizations across the value chain increase energy production, at lower cost, with lower risks, and lower carbon emissions, safely and securely, for generations to come.



The energy systems of the future help put your data to work and provide you with the tools to uncover insights that will optimize energy operations and accelerate innovation. These data-driven solutions are based on:

- ✓ Advanced analytics
- ✓ High-performance computing
- ✓ Artificial intelligence/ Machine learning
- ✓ Next-generation connectivity solutions
- ✓ Internet of Things/ Edge computing

AWS is reinventing the energy industry through innovative cloud technologies and comprehensive digital solutions for:

-  Oil and gas exploration and production
-  The modernized, scalable electrical grid
-  Clean and renewable energy technologies*

*Such as wind, solar, battery storage, clean hydrogen generation and carbon capture and sequestration.

In the chapters ahead, we look at how generative artificial intelligence on AWS will transform the energy industry, and how organizations can get started with harnessing the power of this cutting-edge technology.

BUILT ON AI

Powering Efficiency and Smarter Thinking

Democratizing ML for the energy industry.

Amazon's success has been built on our pioneering work with innovative technologies like AI and ML over 25 years.

Our e-commerce recommendations engine is driven by ML; the paths that optimize robotic picking routes in our fulfillment centers are driven by ML; and our supply chain, forecasting, and capacity planning are informed by ML. Prime Air (our drones) uses deep learning. Alexa, powered by more than 30 different ML systems, helps customers billions of times each week to manage smart homes, shop, get information and entertainment, and more.

AWS has played a key role in democratizing ML and making it accessible to anyone who wants to use it, including more than 100,000 customers of all sizes and industries. Leading energy companies across the value chain have leveraged AWS AI/ML services across several use cases, including preventative maintenance, energy management, demand forecasting, vegetation management and more.



Spain-based [Iberdrola](#), one of the world's largest producers of renewable energy, worked with AWS to harness the power of ML to create a smart assistant that empowers its customers to monitor their own energy use to minimize consumption and reduce their bills.



[Amazon](#) is using a new AI-powered feature that provides a brief paragraph on its product detail pages highlighting a product's features and customer sentiment distilled from written reviews to help customers determine at a glance whether a product is right for them.



[bp](#) is using AWS Sagemaker, which builds, trains, and deploys ML models, to accelerate data science projects and deliver faster data analysis and insights.

And these are just a few of a growing number of [ML energy industry use cases](#).

However, an improved form of artificial intelligence, generative AI, is predicted to deliver an even more profound impact across the energy industry.

THE AI STORY

Changing the Game With Generative AI

Take machine learning to the next level.





We are truly at an exciting inflection point in the widespread adoption of ML, and we believe most customer experiences and applications will be reinvented with generative AI.

Generative AI is a type of AI that can create new content and ideas, including conversations, stories, images, videos, and music. Like most AI, generative AI is powered by ML models—very large models that are pre-trained on vast amounts of data and commonly referred to as foundation models (FMs).

What makes foundation models (FMs) special is that they can perform so many more tasks because they contain such a large number of parameters that make them capable of

learning complex concepts. A parameter is a configuration variable within the model that can be tuned to optimize the model's behavior. While the capabilities and resulting possibilities of a pre-trained FM are amazing, customers get really excited because these generally capable models can also be customized to perform domain-specific functions.

In energy, generative AI has the potential to play a pivotal role in helping companies:

 <p>Increase operational efficiencies</p>	 <p>Reduce health and safety exposure</p>	 <p>Enhance the customer experience</p>	 <p>Minimize emissions associated with energy production</p>
---	---	---	--

...and ultimately accelerate the energy transition



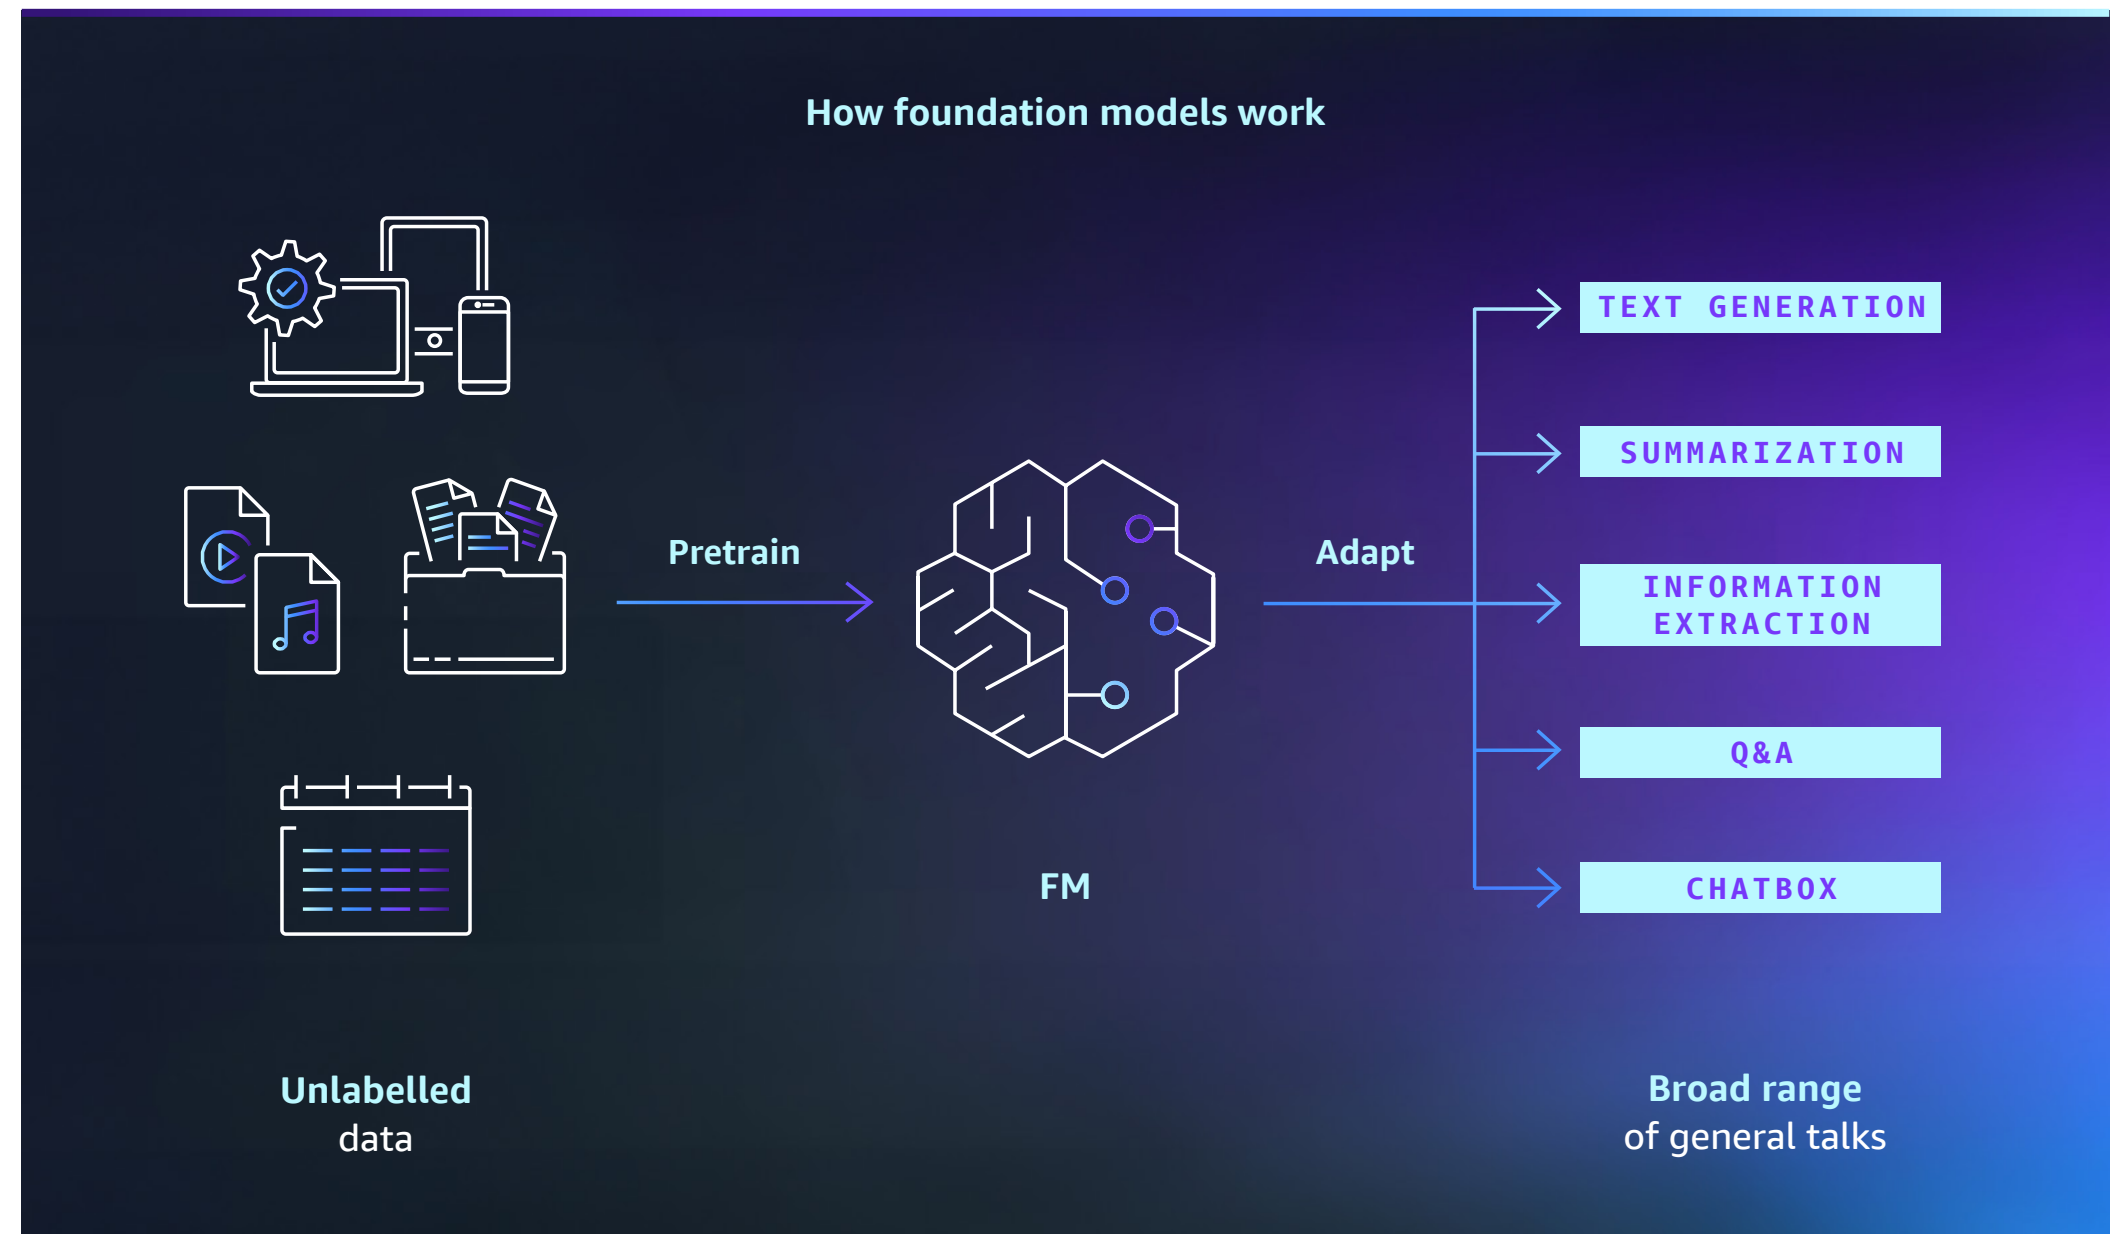
Changing the Game With Generative AI (continued)

How Foundational Models are created

Creating an FM starts with 'pre-training', where a model is trained with terabytes of unlabeled text and/or multi-modal data (such as images, audio, video). This unlabeled data used for pre-training is usually obtained by crawling the Web and contains information from publicly crawled sources (such as Wikipedia and other sites) and proprietary data (if available). During the pre-training process, the model automatically takes context into account from all this unlabeled data, and tracks relationships in sequential data, like the words in this sentence, to develop some understanding of the real world.

A large model with billions of parameters can better capture this knowledge as it is able to store richer and deeper context across large amounts of data in its memory compared to a smaller model trained on a smaller data set. Pre-training models of this size requires access to sufficient quantity and quality of training data, and a large-scale training infrastructure.

Upon completion of pre-training, the resulting FM can deliver performance on a wide range of tasks across multiple domains out of the box. For example, writing blog posts, summarizing documents, solving math problems, engaging in a chat dialogue, answering questions, and even composing poetry.



GENERATIVE AI ENERGY USE CASES

Generative AI in Action

Optimizing safety and sustainability

The challenge

Energy operations often occur in remote, and sometimes hazardous environments that put field workers in danger, while also risking costly loss and damage to vital equipment and other assets.

As a result, the industry has long-sought solutions that help to reduce trips into the field, which would also cut down the carbon emissions from transporting people and equipment.

The solution

Images from cameras stationed at field locations could be sent to a generative AI application that could scan for potential safety risks, such as faulty valves resulting in gas leaks.

The application could generate recommendations for personal protection equipment, along with tools and equipment for remedial work.

This would help remove the need for an initial trip into the field to identify safety issues, minimize the resulting operational downtime, and also reduce the risk to both people and assets.



Reducing cost and downtime

The challenge

Keeping equipment and industrial processes running is vital for fast, efficient energy production.

Any downtime resulting from breakages or faults not only results in costly delays and reduced performance levels but also means replacing expensive assets.


Something which might be avoidable if the problem is spotted before it deteriorates to the point of failure.

The solution

Generative AI could be used to create synthetic data, which is a class of data that is generated, rather than gathered from direct observations of the real world.

In the energy industry, synthetic image data creation could be used for building visualizations of equipment, such as compressors and turbines, with rust, cracks or other issues that could stop them working if not addressed soon enough.

These images could be used to train vision-based ML models to provide predictive maintenance, scanning equipment to deliver advance warning of possible failure. This prevents downtime and helps extend the life of key assets, saving costs.



GENERATIVE AI ENERGY USE CASES

Generative AI in Action

Simulating key energy scenarios

The challenge

One of the most time-consuming and resource intensive areas of oil and gas production is identifying not only where deposits are located, but also evaluating their ease of extraction and the dangers this presents to both people, equipment and the environment.

From a renewables perspective, sourcing suitable geothermal sites throws up similar issues, while planning the location of wind farms is critical to their efficiency and effectiveness. Reducing energy emissions in the atmosphere through carbon capture also presents major challenges in selecting suitable sites.

The solution

Generative AI models could be used to simulate scenarios across the energy spectrum that can save time and resources, while also increasing production efficiency and performance.

Generative Adversarial Networks (GANs) are a popular generative AI technique that could be used to generate synthetic models of oil and gas reservoirs, geothermal sources, carbon capture solutions and wind farm designs and locations based on the wide range of variables in play.

Once the generative AI model is trained, it could be used to generate a large number of synthetic models that could be used for simulation and optimization, improving energy production forecasting.

More efficient data management

The challenge

Accessing and optimizing the value of business-critical data is a constant challenge the energy industry is looking to overcome. This is driven by much of it being decades old, spread across various systems and in different formats.

For example, oil and gas companies have decades of documents created throughout the subsurface workflow in different formats, ranging from PDFs and presentations, to reports, memos, well logs, and Word documents.

This means finding key information can take a considerable amount of time. According to one of the top five operators, engineers spend 60% of their time searching for information.

The solution

Generative AI solutions can work with data from a variety of sources, including intranet pages, bringing it together in a single, common, readable format for easy access and analysis.

This technology could be further augmented through the support of a generative AI-enabled digital assistant, which could be used to help train and assist workers to extract the insight they need, driving faster, better decision making, while optimizing data value.

Contact center transformation

Generative AI smart assistants could also be used by utilities companies to help residential and business customers better understand and manage their energy usage.

The customer experience could be further improved by using generative AI solutions in contact centers to revolutionize call summary and analytics processes for customer-to-agent interactions. The generation of comprehensive summaries and in-depth analysis of recorded and transcribed calls could be automated, empowering call centers to extract valuable insights, enhance service levels, and optimize their operations.




THE PATH TO AI SUCCESS

The AWS Difference

Your route to practical, cost-effective generative AI for energy.

Our combination of expertise and experience in AI and ML, and the energy sector, along with our world-leading cloud platform and cutting-edge tools and applications make AWS uniquely positioned to power and navigate your generative AI journey. Partnering with AWS will give you unrivalled:

**Scalability**

Amazon Bedrock is the easiest way for customers to build and scale generative AI-based applications using FMs, democratizing access for all builders.

**Productivity**

Enable your developers to work more productively and accelerate modelling by generating code suggestions in real-time using AWS Codewhisperer.

**Flexibility**

Choose from a wide selection of open source models from innovative leading startups AI21 Labs, Anthropic and Stability AI, as well as Amazon to find the right one to suit your needs.

**Secure customization**

Customize FMs for your business with just a few labeled examples. Since all data is encrypted and does not leave your Amazon Virtual Private Cloud (VPC), you can trust that your data will remain private and confidential.

**Cost-effective infrastructure**

Get the best price performance for generative AI with infrastructure powered by AWS-designed ML chips and NVIDIA GPUs. Cost-effectively scale infrastructure to train and run FMs containing hundreds of billions of parameters.

**Fast implementation**


Quickly integrate and deploy FMs into your applications and workloads running on AWS using familiar controls and integrations with the depth and breadth of AWS capabilities and services, such as Amazon SageMaker and Amazon S3.

**Responsible AI**

AWS is committed to developing fair and accurate AI and ML services, and provides the tools and guidance needed to build applications responsibly.

Find out more about our innovative tools on the following pages...





Introducing Amazon Bedrock

The easiest way to build and scale generative AI applications.

A new fully managed service from AWS for developing and implementing Generative AI applications, Amazon Bedrock provides access to powerful pre-trained FMs from a range of cutting-edge AI startups and the Amazon-built Titan family of models.

Using Bedrock, energy companies can:




Find out more about Amazon Bedrock and the range of AWS generative AI tools

Room to grow

To optimize generative AI performance and value, it's important to be able to scale applications affordably on a cost-effective cloud infrastructure.

Over the past five years, we have been investing in our own silicon to push the envelope on performance and price for demanding workloads, like ML training and Inference. Our AWS Trainium and AWS inferentia chips offer the best price-performance for training models and running inference in the cloud. Furthermore, our latest NVIDIA chipsets deliver an unmatched GPU-based compute experience, significantly reducing the time to train machine learning models from days to hours.

This ability to maximize performance and control costs by choosing the optimal ML infrastructure is why several of the leading AI startups run on AWS.



AI TOOLS




Transform Coding With Amazon CodeWhisperer

Significantly improve developer productivity.

One area where we foresee the use of generative AI growing rapidly is in coding. Software developers today spend a significant amount of their time writing code that is pretty straightforward and undifferentiated, and trying to keep up with a complex and ever-changing tool and technology landscape. All of this leaves developers less time to develop new, innovative capabilities and services.

Generative AI can take this heavy lifting out of the equation by “writing” much of the undifferentiated code, allowing developers to build faster while freeing them up to focus on the more creative aspects of coding. This is why we have launched Amazon CodeWhisperer, an AI coding companion that uses a FM under the hood to radically improve developer productivity by generating code suggestions in real-time. CodeWhisperer is the only AI coding companion with built-in security scanning (powered by automated reasoning) for finding and suggesting remediations for hard-to-detect vulnerabilities


During the preview of CodeWhisperer, which is now generally available, we ran a productivity challenge, and participants who used the tool completed tasks 57% faster, on average, and were 27% more likely to complete them successfully than those who didn't use the application.



Faster

During a productivity challenge participants who used CodeWhisperer completed tasks 57% faster





Accelerate FM Creation With Amazon SageMaker Jumpstart

Scalable, reliable, and secure model building, training and deployment.

In addition to models in Bedrock, Amazon SageMaker JumpStart is an ML hub offering algorithms, models, and ML solutions. With SageMaker JumpStart, you can access pretrained models, including FMs, to perform tasks like article summarization and image generation.

Pretrained models are fully customizable for use with your data, and you can easily deploy them into production. In addition, you can access prebuilt solutions to solve common use cases, and share ML artifacts, including ML models and notebooks, within your organization to accelerate ML model building and deployment.

For customers who want to create their own FMs, Amazon SageMaker provides managed infrastructure and tools to accelerate scalable, reliable, and secure model building, training, and deployment. Customers can use efficient distributed training to complete training of their own FMs up to 40% faster.

Data security front of mind

To optimize privacy, customer data is neither used to improve Amazon Titan models nor shared with other foundation model providers. It is isolated per customer and remains in the region where it was created. It is also always encrypted in transit with a minimum of TLS1.2 and AES-256 encrypted at rest using AWS KMS managed data encryption keys. Fine-tuned (customized) models are encrypted and stored using customer AWS KMS key. Only customers themselves have access to their customized models.

Generative artificial intelligence

FAIR AND ACCURATE AI

Using AI responsibly

Rising to the ethical challenge.

We are only scratching the surface of how AI applied through ML can transform business and society by tackling some of our most challenging problems, augmenting human performance, and maximizing productivity.

Arguably more than any other technology to date, AI raises important ethical issues. This makes its responsible use key to fostering continued innovation. AWS is committed to developing fair and accurate AI and ML services, and providing the tools and guidance needed to build applications responsibly.

Our range of responsible AI resources include:



The Responsible Use of Machine Learning Guide

Provides considerations and recommendations for responsibly developing and using ML systems across design and development, deployment ongoing use.



Expert advice

Work with experts in responsible ML to create an operational approach encompassing people, processes, and technology that maximizes benefit and minimizes risk.



Education and training

Continuous education on the latest developments in ML and its responsible use through programs like the AWS Machine Learning University (Bias and Fairness Course), Training and Certification program, and AWS ML Embark.

Prediction monitoring tools

A range of services that help you:



Detect bias in datasets and models

Amazon SageMaker Clarify helps you mitigate imbalances in datasets by detecting potential bias during data preparation, after model training, and in your deployed model by examining specific attributes.



Better understand model predictions

Amazon SageMaker Clarify provides greater visibility into model behavior, both overall and for individual predictions, providing transparency for stakeholders, more deeply informing humans making decisions, and tracking whether a model is performing as intended.



Monitoring and human review

Amazon SageMaker Model Monitor detects inaccurate predictions from models deployed in production. Amazon Augmented AI enables human reviews of ML predictions when human oversight is needed.

INVESTING IN THE FUTURE

The AWS Generative AI Innovation Center

Accelerate the development of more effective AI concepts and solutions.

To help energy companies better build and deploy their own generative AI products and services, we've invested \$100 million in launching the AWS Generative AI Innovation Center.

The initiative will connect our AI and ML experts with energy businesses around the globe to help envision, design and turn new generative AI concepts into reality.

The center's team of strategists, data scientists, engineers, and solutions architects, as well as experts from the AWS Partner Network, will work step-by-step with businesses to build bespoke solutions that harness generative AI in the most effective and responsible way.

Our generative AI experts will work closely with the energy sector to develop use cases that will help overcome key industry challenges, from detecting resources and more efficient production to accelerating the transition to renewables and improving the customer experience.

Working with AI, together – how we can collaborate to accelerate:

Education

We will conduct sessions to provide context around generative AI, discuss the art of the possible, review our offerings, and learn more about your short and long-term needs.

Workshop

We will work with you to identify the key challenges that generative AI can help you overcome as a business.

Engagement

We will showcase the power and expanse of generative AI on AWS against a targeted use case and presenting the results to key stakeholders.

Production plan

We will provide an Execution Roadmap and support you in developing production-ready generative AI use cases, including considerations for scale.



Get Started with Generative AI today

Discuss your energy challenges with our AWS experts, who will work with you towards a proof of concept and guide you along your generative AI journey.

We'll align business and technical stakeholders, creating a roadmap that will lead you to the right solutions to transform your business and help you meet your key goals.

Let's set up a meeting now.