




How technology leaders transform mission impact with generative AI in AWS GovCloud (US)



CONTENT

Overview
A leader’s guide to using your data as a differentiator **3**

Generative AI
AWS is a long-term AI innovator in highly regulated cloud **4**
Generative AI is ushering in a new technology era **5**

Section 1
What’s possible with AI? **6**

Section 2
What should I do first? **8**

Section 3
AWS GovCloud (US) cloud technology
underpins the generative AI journey **10**

Conclusion
The bottom line on generative AI
and regulated cloud solutions **11**



OVERVIEW

How technology leaders transform mission impact with generative AI in AWS GovCloud (US)

Technology leaders in the public sector want to leverage generative AI to promote efficiency, enhance mission impact, and achieve business goals. At the same time, they are obligated to safeguard sensitive government data through alignment with government regulatory compliance standards.

How can technology leaders efficiently deliver generative AI innovation while aligning with the broad spectrum of strict regulatory requirements? The answer is simple: [Amazon Web Services \(AWS\) GovCloud \(US\)](#). AWS provides the technology that underpins a solid foundation for securely and compliantly building and deploying generative AI capabilities.

AWS is a long-term AI innovator in highly regulated cloud

More than a decade ago, AWS realized that technology leaders with U.S. data residency, U.S. sovereignty, EAR/ITAR, FedRAMP High, and/or DoD SRG IL-4/5 compliance requirements need secure and compliant access to AI capabilities with the supporting cloud infrastructure to promote large-scale innovation. We learned that AWS GovCloud (US) customers need to deploy AI solutions that improve citizen experiences, boost employee productivity, and accelerate process optimization in areas including classifying images, enterprise searches, processing and analyzing documents, transcribing audio and video, translating languages, and conversation chat. AWS answered this need many years ago by deploying a suite of compliant AI capabilities, including [Amazon SageMaker](#), [Amazon Rekognition](#), [Amazon Kendra](#), [Amazon Textract](#), [Amazon Transcribe](#), [Amazon Translate](#), [Amazon Polly](#), and [Amazon Lex](#).



Generative AI is ushering in a new technology era

Many technology leaders operating highly regulated, on-premises, and cloud environments began in 2025 with vast opportunities to innovate with generative AI in the cloud efficiently. This is because they realized this new generative AI technology era provides innovative possibilities to elevate business and mission outcomes in ways that previously seemed impossible or uneconomical to achieve.

This is why AWS introduced managed generative AI capabilities to AWS GovCloud (US) with the availability of [Amazon Bedrock](#) in 2023, a managed generative AI service to build and scale applications using foundation models. Amazon Bedrock is available with [Anthropic Claude Sonnet](#) and [Haiku](#) foundation models, [Meta Llama](#) foundation models, and [Amazon NOVA family](#) foundation models. [Amazon SageMaker Jumpstart](#) enables the use of publicly available, open-weight foundation models sourced from Hugging Face and Meta. Bookmark our [What's New with AWS GovCloud \(US\)](#) page for more updates.

As part of Amazon's commitment to transformative AI and to promote efficient innovation for the government and the supporting industrial base, AWS entered into a four billion dollar strategic partnership with Anthropic, one of the world's leading foundation model providers and a leading advocate for generative AI innovation.

The value of AWS generative AI investments for technology leaders operating in highly regulated environments is profound: AWS GovCloud (US) enables a fusion of innovation and efficiency by providing a secure, compliant, sovereign, and isolated cloud foundation for generative AI that can rapidly transform government service delivery and operational efficiency.

These cloud-based generative AI capabilities facilitate workforce productivity gains by automating routine documentation, analysis, and administrative tasks, allowing staff to focus on high-value activities. Real-time data analysis capabilities accelerate operational mission decisions by quickly processing vast amounts of information into actionable insights. Operations may be streamlined through AI-enhanced workflows that reduce processing times and human error while maintaining detailed audit trails. Research and discovery are accelerated through AI-assisted analysis, pattern recognition, and predictive analytics that enable scientists to create greater value from their data while optimizing resource allocation and mission outcomes.

What's possible with AI?

AWS AI capabilities, infrastructure, and compliance programs position technology leaders to efficiently deploy AI capabilities using AWS GovCloud (US). Here are some examples:

Transform and enhance citizen experience




AI-enabled [Amazon Lex](#) chatbots in webpages and contact centers, powered by [Amazon Connect](#), can increase customer satisfaction and call velocity. Conversation analytics during customer interactions may be used to reduce time to resolution, improve customer experience, and create tailored customer experiences.

Elevate patient care quality and outcomes



AI-enabled [Amazon Textract](#) and [Amazon Comprehend](#) underpin [intelligent document processing \(IDP\)](#) solutions that not only extract text and structured data from documents but also employ AI capabilities to provide business insights and determine relationships between patient datasets. The velocity of IDP helps healthcare providers efficiently manage patient care and automatically reduce the potential for information errors, so providers can focus on positive patient health outcomes rather than on managing patient records and bills.




SECTION 1

Improve employee experience and impact



AI-enabled [Amazon Kendra](#) provides employers with solutions to help employees spend less time searching intranets and data repositories or opening countless documents to find answers to their questions. This results in increased productivity and decreased stress. AI-enabled [Amazon Transcribe](#) and [Amazon Comprehend](#) enable employees to [summarize meetings](#) and dialogues through automated audio note collection, which saves employees time and improves information accuracy. This empowers leaders to make more informed business decisions and expedite decision-making.

Elevated and enhanced scientific data analysis



AI enabled the Energy Department, in [partnership](#) with the Treasury Department, to apply advanced analytics to complex data structures, resulting in a deeper understanding of crypto financial data and promoting scientific discovery.

SaaS and ISV offerings



The regulated software as a service (SaaS) and independent software vendor (ISV) industrial base can also take advantage of the innovative possibilities AI technology offers to enable sustainable growth, expand market presence, strengthen brand quality, modernize legacy systems, increase security, drive efficiencies, and reduce environmental impact.

These are just a few examples of how technology leaders can take advantage of the AI capabilities available in AWS GovCloud (US). Outlined below are ways technology leaders can start and efficiently promote their AI journey in AWS GovCloud (US).

What should I do first?

What can technology leaders do to efficiently deploy generative AI in this new technology area? Here are three pieces of advice from the AWS GovCloud (US) team:

1 Start experimenting today

Start with a small yet well-defined generative AI project with short-term measurable outcomes, broad executive support, and minimal risk. Demonstrate the success of this first project, and then kick off a slightly larger project with more impactful measures. Keep repeating this cycle in time-bound iterations and continue to demonstrate increasing incremental mission impact. This is how Amazon started over two decades ago when we launched our [Amazon.com](https://www.amazon.com) e-commerce recommendations engine, underpinned with AI technology. Technology leaders will benefit more from this approach than endless planning and waiting for the “perfect time” to start their generative AI journey.

2 Implement a data strategy that supports generative AI innovation

Creating a robust enterprise data strategy is critical to implementing successful generative AI solutions. Generative AI and foundation models have elevated the need for clean training data. Quality and clean data are necessary to fine-tune generative AI models to unlock and maximize business and mission value. Technology leaders sometimes face situations where their data strategy does not provide quality training data, which causes the results from generative AI solutions to produce inferior outcomes due to the “garbage in, garbage out” phenomena. In fact, “Training models with poor data quality will lead to poor results,” says Werner Vogels, VP & CTO of Amazon.com. “You will need to filter out bias, hate speech, and toxicity. You’ll need to ensure that the data is free of personally identifiable information (PII) or sensitive data, and make sure your data is deduplicated, balanced, and doesn’t lead to oversampling.”

SECTION 2

A well-defined data strategy will pay dividends to technology leaders on their generative AI journey. Whether you are building your model or customizing one, all leaders need a data strategy that ensures relevant, high-quality data is available. Some data may be more than 20 years old residing on a mainframe, while others might be massive, unstructured datasets living in a legacy storage system. Data must be up-to-date, complete, accurate, discoverable, and available when needed. AWS can help accomplish these goals with [Amazon Redshift](#) for data warehouses, [Amazon Simple Storage Service \(Amazon S3\)](#) for data lakes, and [Amazon EMR](#) for big data. In turn, these services can increase the value of data as a component of an enterprise generative AI solution.

3 Partner with AWS to promote positive, mission-driven outcomes

AWS offers customers access to generative AI-focused, AWS GovCloud (US) specialist solutions architects and professional services. AWS also has a broad partner community of consultants and integrators. What should you do if you are not sure how to use generative AI to enhance mission impact and increase government efficiency? Just visit [AWS in the Public Sector](#) website and fill out our contact form to request a call. You can also visit the [AWS AI solutions library](#) to learn more about generative AI use cases and reference architectures.

A well-defined data strategy will pay dividends to technology leaders on their generative AI journey.

AWS GovCloud (US) cloud technology underpins the generative AI journey

Providing a performant, secure, and compliant technology platform to host generative AI workloads in an on-premises environment can be prohibitively costly for a typical technology leader. With AWS, you pay as you go. The benefits of AWS GovCloud (US) as a platform for generative AI solutions are extensive. AWS GovCloud (US) supports critical sovereignty, security, and compliance programs that support generative AI workloads. AWS GovCloud (US) regions are inside an ITAR, FedRAMP High, and DoD SRG IL-4/5 compliance boundary. AWS GovCloud (US) cloud services are authorized at FedRAMP High and DoD SRG IL4 and IL5 or are on a path to authorization. This compliance posture provides customers with a trusted, isolated, and proven path to receive the authority to operate. The simple, pay-as-you-go, cloud consumption model enables customers to consume generative AI services, database services, GPUs, CPUs, storage, and networking for specified amounts of time in an on-demand, compliant, and scalable model. No longer are there requirements to build and deploy generative AI, on-premises platforms that will sit idle for intermittent periods, awaiting consumption or for users to wait for resources to become available due to resource scheduling constraints.

AWS GovCloud (US) comprises two isolated and U.S. Sovereign Regions that promote resilient architectures. While they are available on the public internet, they are isolated from the AWS global footprint and operated independently by U.S. citizens. Each Region contains three Availability Zones composed of data centers with redundant power, networking, and connectivity. Availability Zones are located far enough apart to reduce the risk of a single event impacting availability, yet near enough to enable synchronous replication, rapid failover, and low latency. This Region design helps protect applications against disruptions, such as human mistakes, unexpected traffic spikes, utility failures, earthquakes, and weather events. The resulting business and mission value is that technology leaders can architect for the level of availability and resiliency their generative AI application requires and their budget allows.

CONCLUSION

The bottom line on generative AI and regulated cloud solutions

The intersection of cloud computing, services, networking, and data storage with security and compliance, enables highly regulated customers to efficiently deploy generative AI solutions on AWS GovCloud (US). By embracing its agility and scalability, technology leaders can accelerate mission outcomes to meet evolving government technology needs in the modern digital age, creating a more responsive and future-ready compliant technology platform for the nation's most critical technology assets. AWS encourages technology leaders to think big, experiment, innovate, and take advantage of AWS GovCloud (US) as a strategic technology enabler to achieve their generative AI-based mission outcomes.



ABOUT THE AUTHORS



David Schatzman

David is a technical business development manager for Amazon Web Services (AWS), focused on serving public sector civilian and financial customers using the AWS GovCloud (US) Regions. In this role, David works closely with customers to ensure alignment of their mission goals and technology strategies with the capabilities of the AWS GovCloud (US) Regions. David is also interested in global economics, digital assets, cloud security, and cloud resiliency and is the lead for the AWS GovCloud (US) digital assets, supercomputing, perimeter protection, modernization, and resiliency product strategies. He is a doctoral candidate at the Liberty University School of Business and holds an MBA, MS, and BS, as well as several Project Management Institute credentials, such as the PfMP.



Michael Greenwald

Michael is the global head of financial Innovation and digital assets at Amazon Web Services (AWS) and leads AWS Global Executive Relations. He works with US and international governments on cloud computing and is responsible for emerging technology innovation and implementation. In 2023, he was appointed to represent Amazon on the US Commodity Futures Trading Commission's (CFTC) Technology Advisory Committee. He previously served as the first US Treasury attaché to Qatar and Kuwait and has served in senior roles with two presidential administrations and under three Treasury Secretaries.



Scott Bourn

Scott is a technical business development manager for Amazon Web Services (AWS), driving ISV and SaaS partner success in the public sector using the AWS GovCloud (US) regions. In this role, Scott works closely with partners to break through market barriers and drive campaign fulfillment that is aligned with their strategic plan and goals, all with the capabilities of the AWS GovCloud (US) regions. Scott is also interested in privacy, cyber security, network virtualization, cloud FinOps, edge computing, and cloud resiliency and is the lead for the AWS GovCloud (US) telecommunication and network provider partnerships. Prior to AWS, Scott served as a Senior Manager for Deloitte Consulting's public sector-focused cloud-managed services and as Program Manager for Akamai Technologies' DoD-managed CDN. Scott earned an MBA at the University of Maryland and holds PMP, CISSP, ITIL, and DevOps certifications.



Shawn Asfeld

Shawn is a Senior Solutions Architect for AWS GovCloud (US). He has extensive experience working with federal and civilian customers to build a wide variety of secure and compliant workloads in AWS GovCloud (US). Shawn's current focus is helping customers and partners build solutions on AWS GovCloud (US) to meet various levels of compliance, including FedRAMP, CMMC, and DoD authorization. He earned a BS degree from Texas A&M University and is an AWS-certified solutions architect professional.



Aaron Sengstacken

Aaron is a machine learning specialist solutions architect at Amazon Web Services. Aaron works closely with public sector customers of all sizes to develop and deploy production machine learning and generative AI applications. He is interested in machine learning, technology, and space exploration. Aaron earned a BS in mechanical engineering from the University of Missouri and an MS in aeronautics and astronautics from Purdue University.