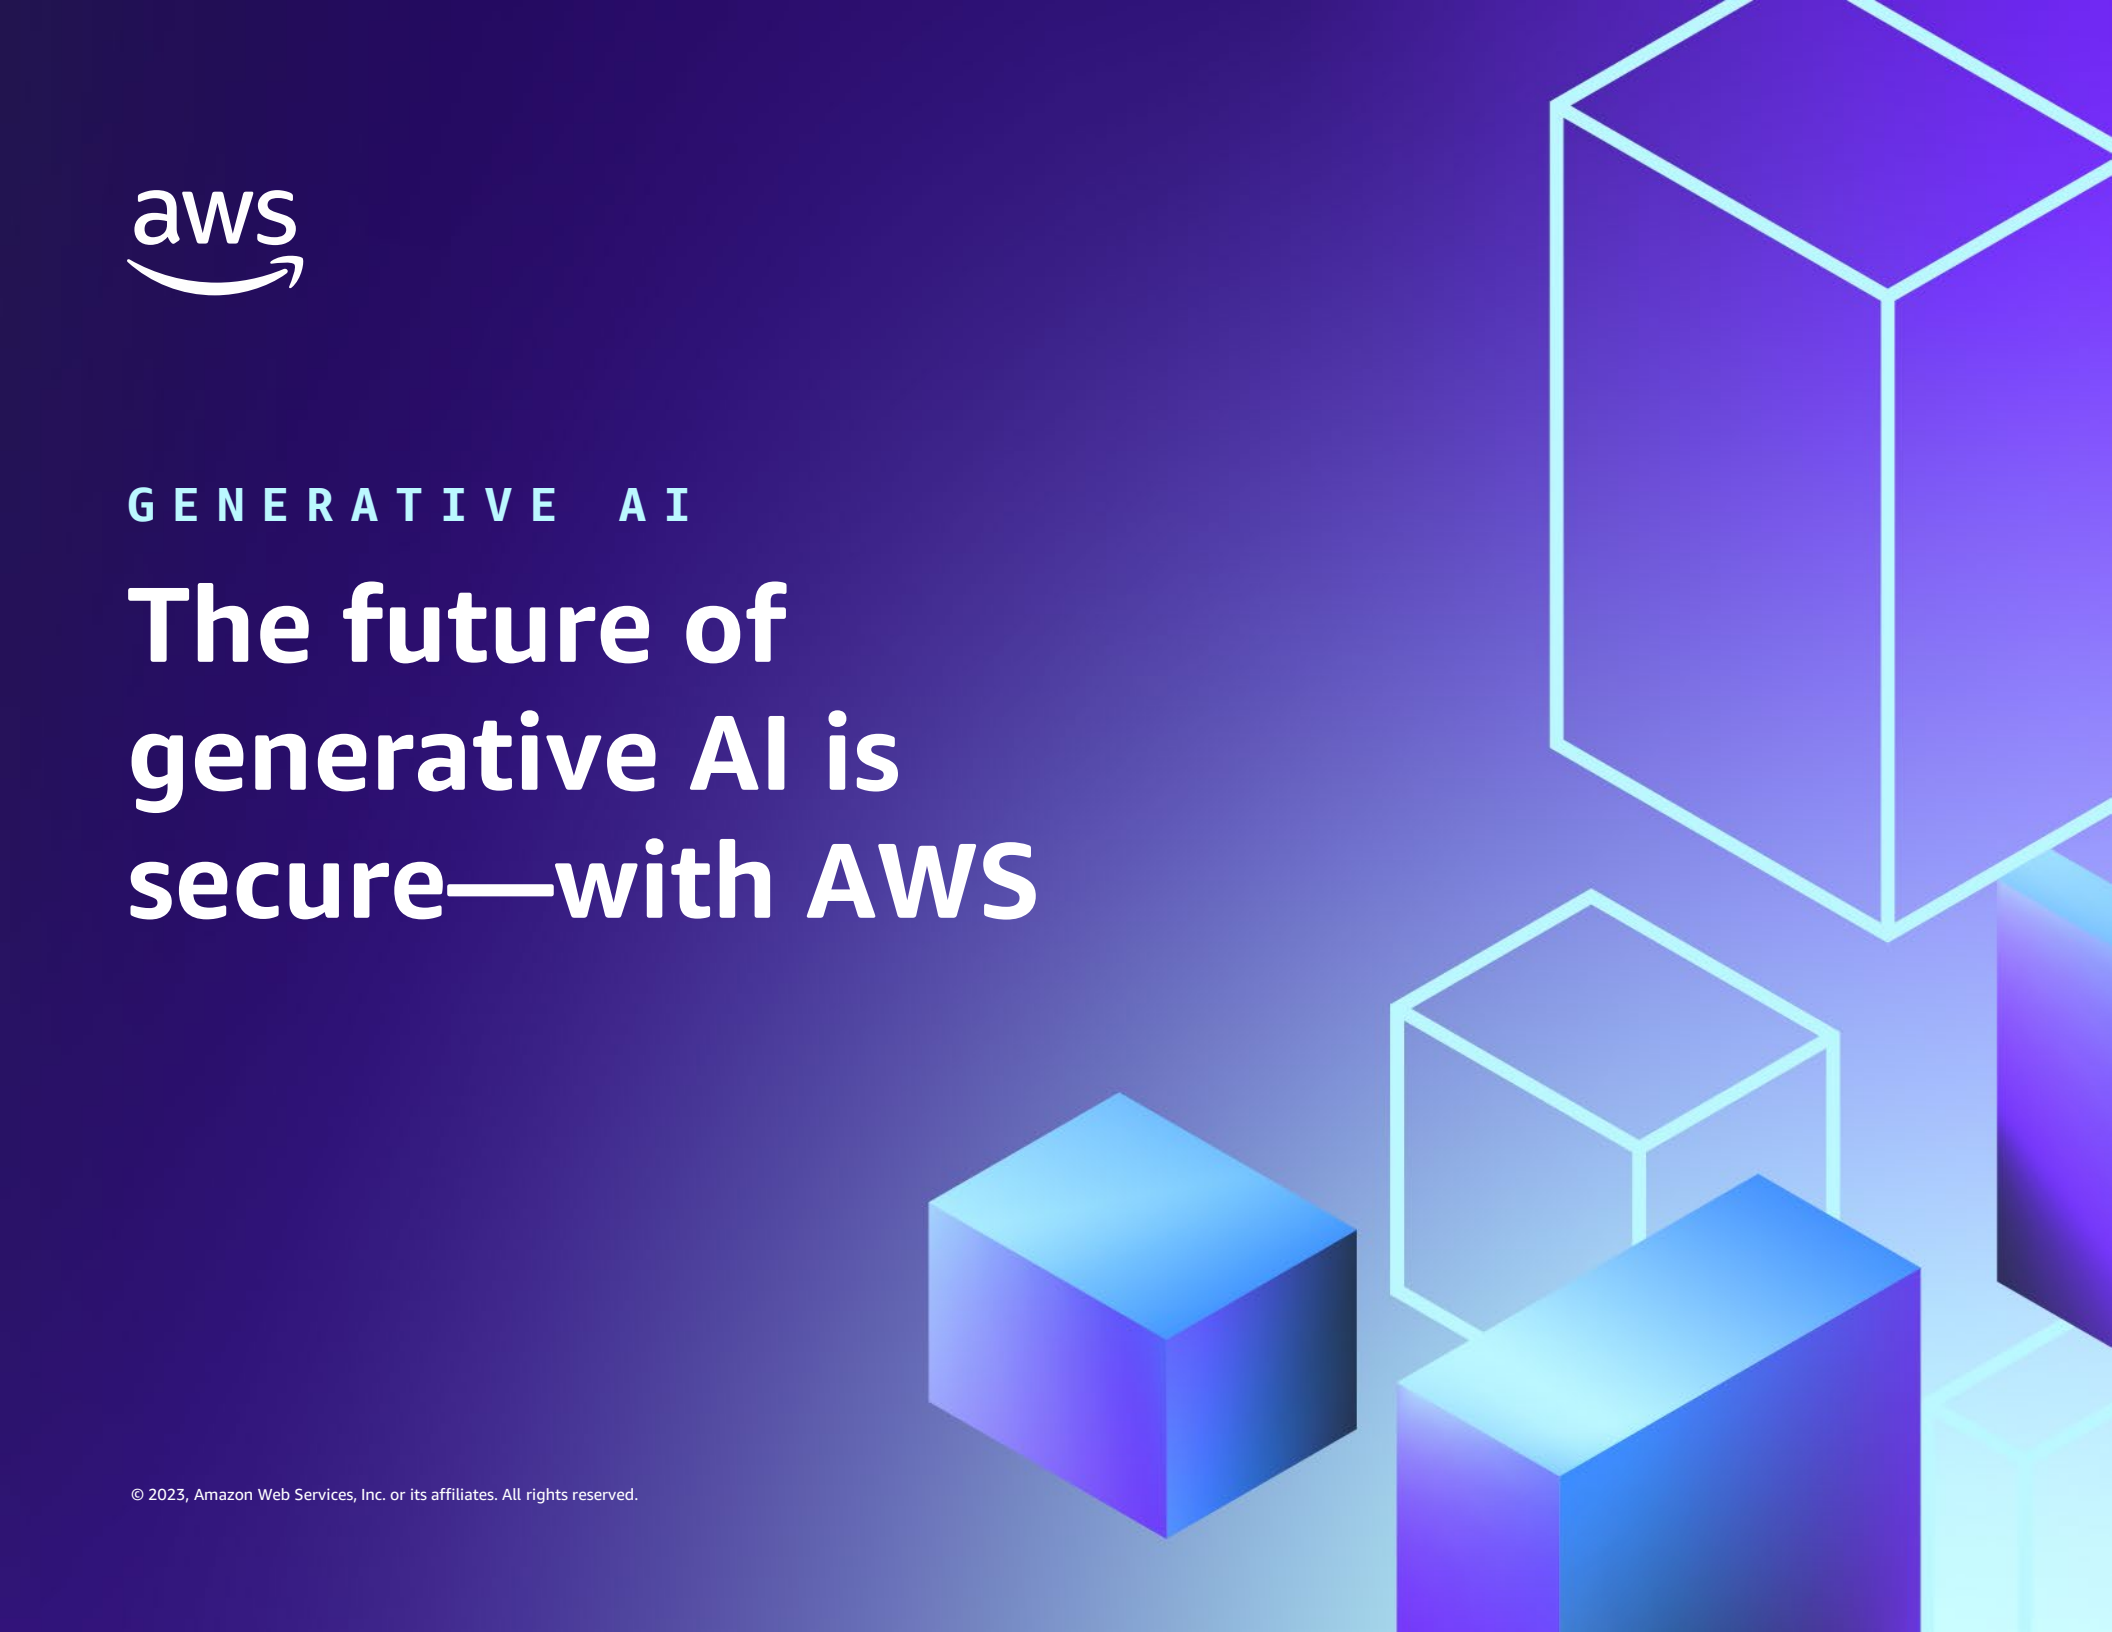

GENERATIVE AI

The future of generative AI is secure—with AWS

Table of contents

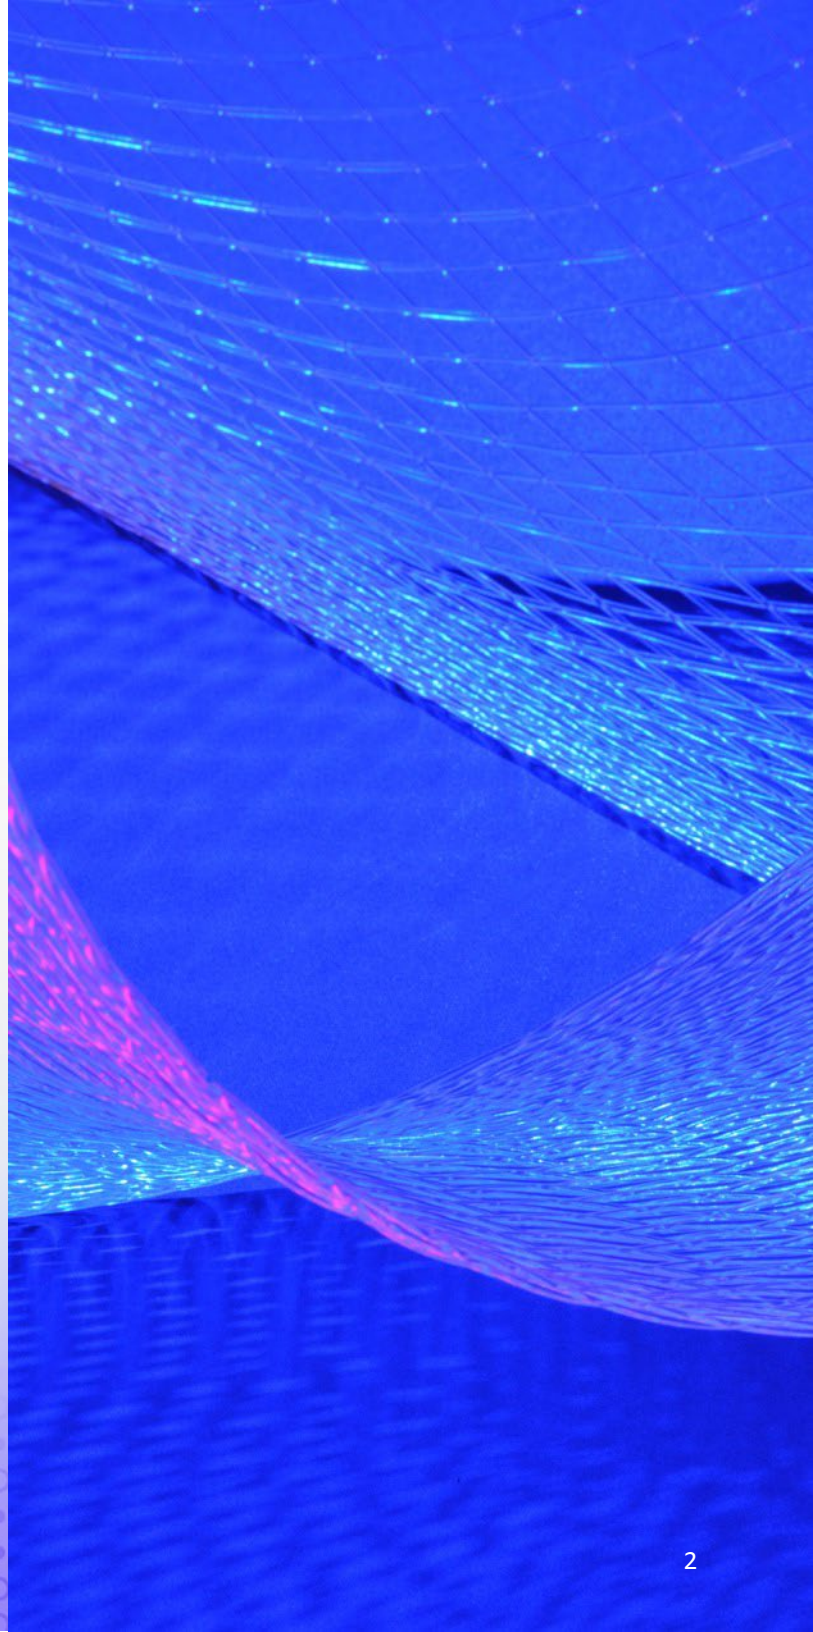

When AI leads, security concerns follow

Generative AI holds big promise for the enterprise. Developed to produce text, images, and other media, generative AI is democratizing content creation, summarization, and document processing. Services like ChatGPT and DALL-E are simple yet extremely powerful tools used within organizations around the world. And as generative AI is added to proprietary solutions, organizations are realizing tangible benefits across virtually all business units. In IT security, for example, generative AI is augmenting threat detection, adversarial defense, and network security.

56% of workers are already using generative AI on the job (Source)

However, without clear security guidance and governance, generative AI is also raising security and privacy concerns. Employees are inputting incredibly valuable IP into these services. For security teams, it is imperative that any information shared externally stays protected, secure, and private. But data privacy is only one factor. Model bias, the creation of harmful content such as deepfakes, and the poisoning of models through malicious input are other reasons to approach generative AI with care.

Looking ahead, organizations must develop a robust and effective artificial intelligence (AI) security strategy including:

1. Securing each use
2. Compliance
3. Resilience consideration

This ebook examines the three key areas of a generative AI security strategy.

- Securing generative AI
- Securing against generative AI
- Developing generative AI for security

When designing your generative AI security strategy, ask:

- What do I need to protect?
- What industry and governance standards am I beholden to?
- What needs to be visible to organizational leaders?
- Who are the intended users and what are the guidelines for appropriate use?

Infrastructure, services, and partners to elevate your cloud security

As you design your generative AI security strategy, keep in mind, security is the top priority for Amazon Web Services (AWS). AWS offers global cloud infrastructure architected to be the most secure of any public cloud. AWS has more than one million active users—including the most security-sensitive organizations like government, healthcare, and financial services—building, migrating, and managing applications and workloads on the cloud. Plus, the Shared Responsibility Model makes it easy to understand your choices for protecting your unique AWS environment, and it provides access to resources that can help you implement end-to-end security quickly and easily.

- Detection and response
- Data protection
- Network and application protection
- Compliance

Security OFF the cloud model

Securing generative AI

- Amazon Bedrock
- Amazon SageMaker JumpStart

Securing against generative AI

- Amazon Macie
- AWS Network Firewall

Security IN the cloud model

Developing generative AI for security

- Amazon CodeWhisperer
- Amazon Detective Finding Groups
- Amazon GuardDuty

Reduce risk with automated security services

With AWS services you can automate security tasks—reducing human configuration errors and giving your team more time to focus on critical work. AWS has a wide variety of integrated solutions that can automate tasks, make it easier for security teams to work with developers and operations teams, and deploy code faster and more securely. For example, automating infrastructure and application security checks allows you to continually enforce your security and compliance controls and help ensure confidentiality, integrity, and availability at all times.

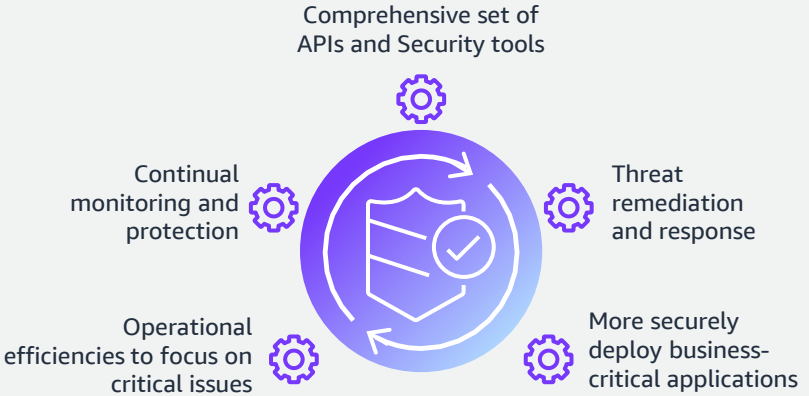

Tap into a broad ecosystem of security competency partners

In addition to security services, you can also extend the benefits of AWS by using security technology and consulting services from AWS Professional Services and the AWS Partner Network. AWS has carefully selected providers, many of whom specialize in delivering security-focused solutions and services for your specific workloads and use cases. Easily find, buy, deploy, and manage cloud-ready software solutions in a matter of minutes from AWS Marketplace.

Securing generative AI

Generative AI applications are powered by foundation models (FMs) that are trained on vast quantities of data. FMs analyze this data to identify patterns and learn how to generate new, similar content. To build generative AI applications that meet your specific business requirements, you will typically need to customize an existing FM by training it on your organization's data.

This data may include proprietary information, valuable intellectual property, and sensitive information about your customers, so ensuring its security is critical. Take into consideration these steps to ensure the safety and privacy of generative AI applications.

Customize FMs for your business with just a few labeled examples. **None of your data is used to train the underlying models.** Since all data is encrypted and does not leave your VPC, you can trust that your data will remain private and confidential.

Keep your data private

With Amazon Bedrock you can build your own generative AI application. Amazon Bedrock is a fully managed service that makes FMs available through an API. Using this service, you can customize FMs privately and bring in your own data. Through an API endpoint you can access Amazon Bedrock either through your public address space or internet from your corporate network using a NAT gateway.

Keep in mind, the traffic never goes over the internet. It goes over the same address space in the same region and it never exits your private network or network border. In addition, all traffic is encrypted and never leaves your virtual private cloud.

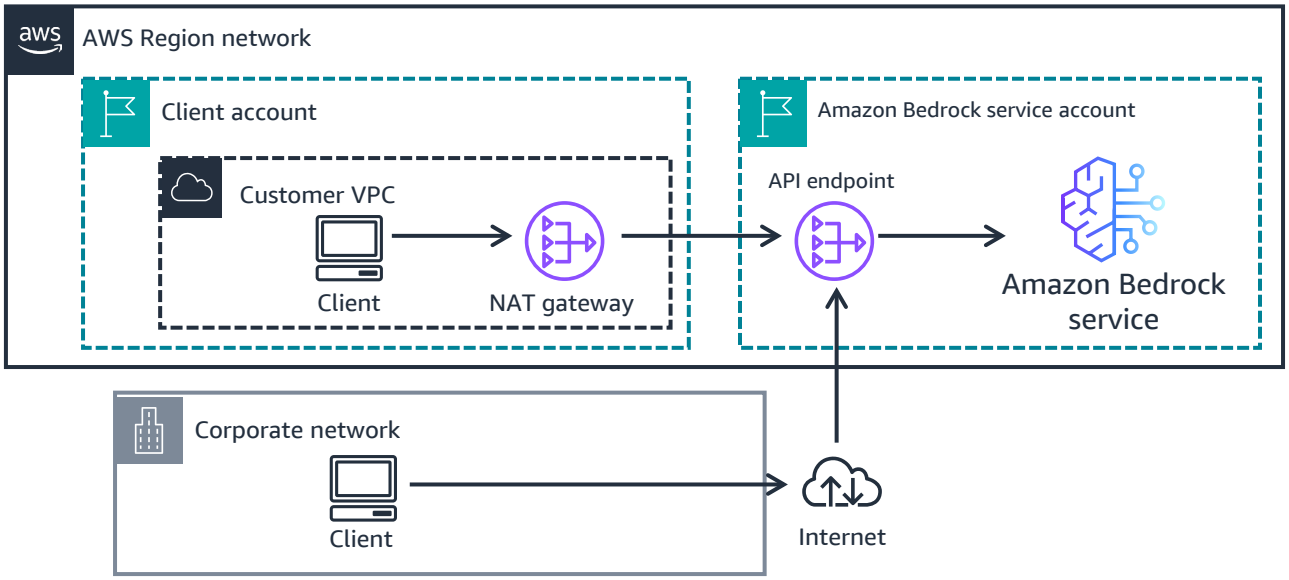

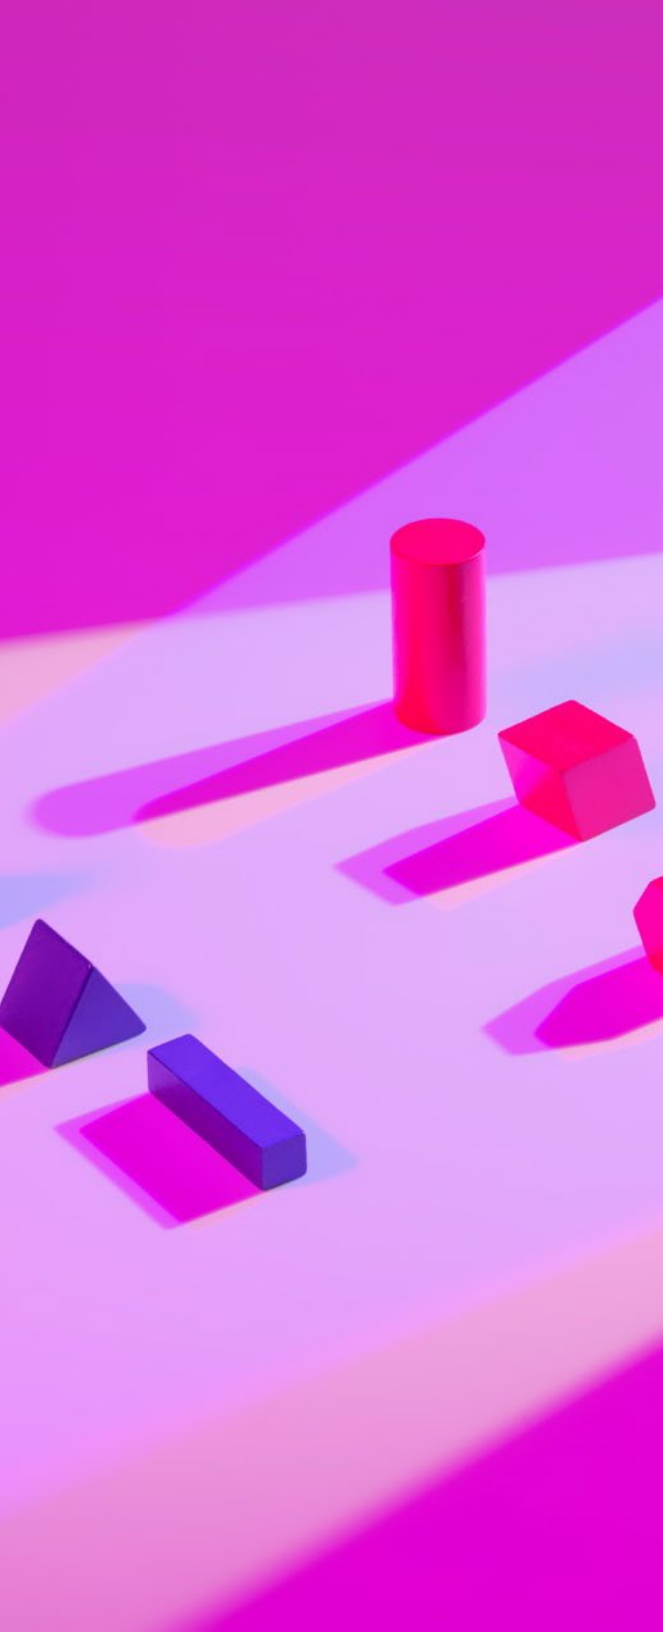

Amazon Bedrock provides bar-raising security controls. You get the standard AWS Identity Access Management (AWS IAM) controls for authentication and the ability to continuously monitor, log, and retain account activity with AWS Config and AWS CloudTrail. All your data is encrypted at rest using your own AWS Key Management Service (Amazon KMS) keys, which provides full control and visibility into how your data and custom models are being stored and accessed.

Amazon Bedrock can also attach its training instances to your Amazon Virtual Private Cloud (Amazon VPC) in order to read from and write to Amazon Simple Storage Service (Amazon S3). And, if you set up a single tenant in Amazon Bedrock, the service can attach its inference instances to your Amazon VPC to read from and write to Amazon S3.

Model tenancy

Single-tenant endpoint	Single-tenant endpoint
1. Deployment available to a single customer 2. Holds a single version of a baseline 1P/3P model that has been fine-tuned by a customer	1. Deployment available all customers 2. Holds a baseline version of each supported 1P/3P model

3. No inference request's input or output text is used to train any model(s) in the deployment
4. Model deployments are inside an AWS account owned and operated by the Bedrock service team
5. Model vendors have no access to any customer data

Fine-tune ML models

To secure generative AI at the application level, you must continuously identify, classify, remediate, and mitigate any vulnerabilities in inputs, outputs, and the model itself. Using Amazon SageMaker JumpStart, you can easily deploy and fine-tune natural language processing models to help your organization meet the strict security requirements of machine learning workloads.

Quickly integrate and deploy FMs into your applications and workloads running on AWS. **Use familiar controls and integrations** with the depth and breadth of AWS capabilities and services like Amazon SageMaker and Amazon S3.

Securing against generative AI

As with any other tool, generative AI can introduce the potential for misuse. There have already been examples of using generative AI for phishing emails, social engineering attacks, and other types of malicious content. As threat actors begin to abuse the technology, AWS is preparing for opportunities and challenges that lie ahead.

However, while generative AI changes how code is created, it does not change how the code works. Certain attacks may be simpler to deploy, and therefore more numerous, but the foundation of how AWS detects and responds to these events remains the same.

Deploy end-to-end security with AWS services

When you build on AWS, you have native cloud services at your disposal to create end-to-end security—from identifying risks to remediation. AWS also offers guidance to help you strengthen your security posture at every step of the way, ensuring your organization is protected from cyberattacks.

Discover and protect sensitive data with Amazon Macie

Amazon Macie is a data security service that uses ML and pattern matching so you can discover sensitive data—before it becomes the target of an attack.

Protect traffic across your VPCs with AWS Network Firewall and security groups

With AWS Network Firewall, you can define firewall rules that provide fine-grained control over network traffic. Use security groups in Amazon VPCs to further control traffic that's allowed to reach and leave your AWS resources. This allows you to protect your managed infrastructure and secure it against attacks that use generative AI.

aws

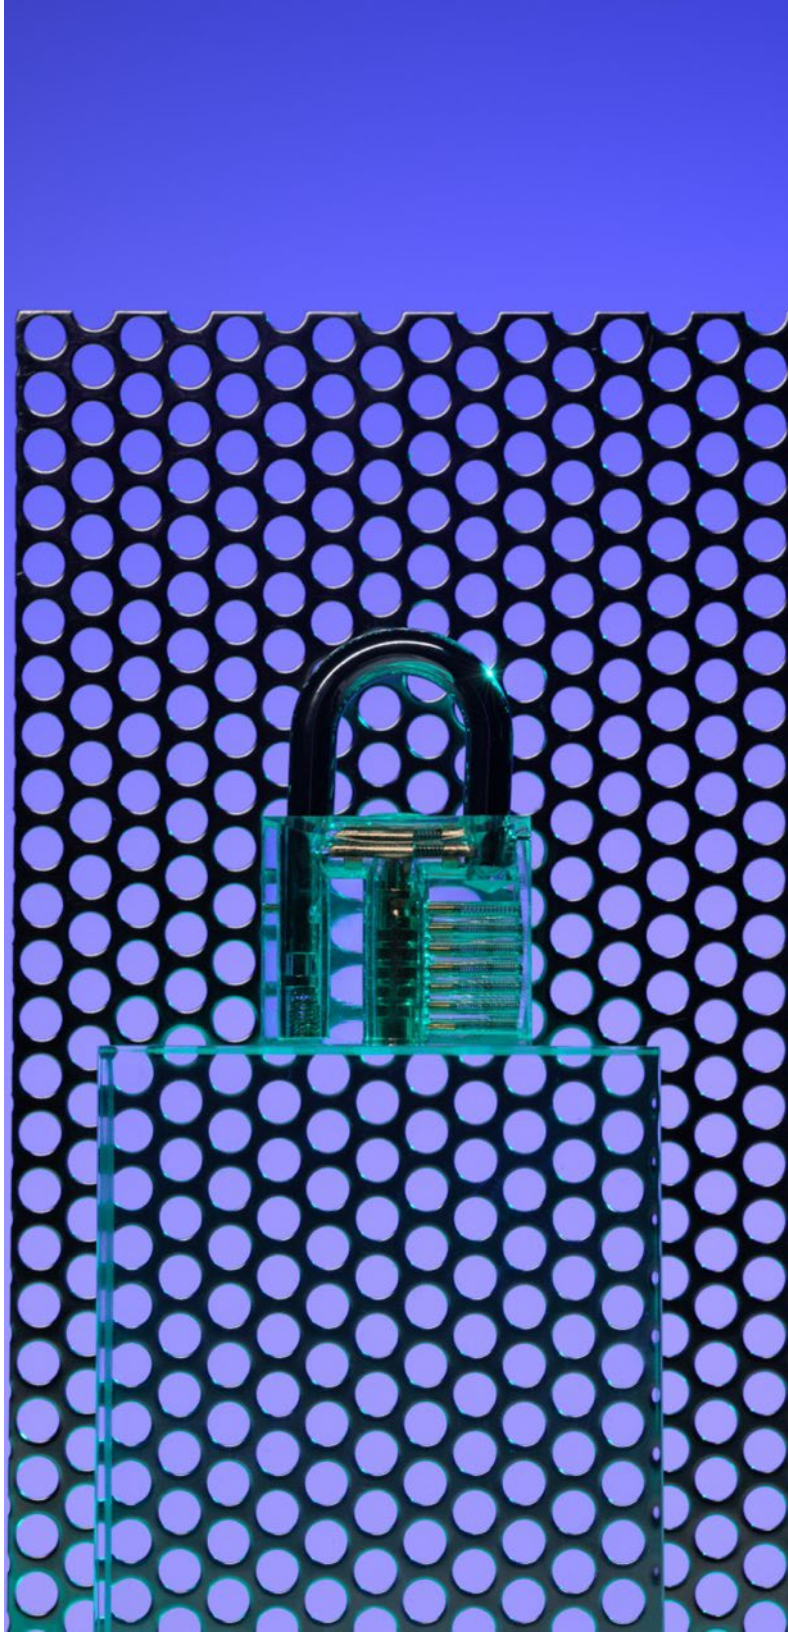

Generative AI for security

Beyond securing your generative AI applications and keeping data private, generative AI can also be used as an indispensable tool for security engineers. From AI-generated security fixes, to assessing vulnerabilities in IAM configurations, generative AI and large language models (LLMs) can free up security teams to focus their energy on more strategic business initiatives.

Use cases for generative AI in security

- AI-generated security fixes
- Weekly reviews of IAM configurations
- Prioritized security alerts
- Automated responses to threats
- Deeper investigation of logs
- Error reports via chatbots

Generate code suggestions for more secure builds

Recognizing this potential, AWS is continuing to invest in generative AI solutions. Services such as Amazon CodeWhisperer have generative AI built in to help you write more secure code and improve developer productivity. As an AI-powered code companion, you can generate code suggestions in real time for AWS services such as Amazon Elastic Compute Cloud (Amazon EC2), AWS Lambda, and Amazon Simple Storage Service (Amazon S3). With security scans that can be run in the IDE, potential vulnerabilities can be found and corrected earlier in the application lifecycle—lowering the cost, time, and risk of application development. Amazon CodeWhisperer is an AI coding companion with built-in security scanning on code for hard-to-detect vulnerabilities.

Trained on billions of lines of Amazon and open source code, **Amazon CodeWhisperer** is an AI coding companion that helps you quickly write secure code. It generates whole line and full function code suggestions in your IDE in real time, based on your natural language comments and surrounding code.

Deploy services that are powered by AI

A feature in Amazon Detective called Finding Groups uses machine learning to distill thousands of security findings from connected security events. This makes it easier for security analysts to understand the complex interactions that result from a potential issue or security event. Finding Groups works by analyzing thousands of unique security findings aggregated from AWS Security Hub across hundreds of AWS resources.

Amazon GuardDuty offers intelligent threat detection. Using machine learning and anomaly detection, Amazon GuardDuty identifies previously difficult-to-find threats, such as unusual API call patterns or malicious AWS IAM user behavior. Amazon GuardDuty also has integrated threat intelligence, which includes lists of malicious domains or IP addresses from AWS security and industry-leading third-party security partners.

Get started with smarter, safer generative AI

The possibilities of generative AI are just opening up. For security teams, now is the time to make sure your organization adopts this indispensable technology in a way that's safe, secure, and beneficial. As you build out your security strategy, consider how to secure generative AI applications, secure against generative AI, and develop generative AI for security.

AWS is committed to helping you build generative AI applications that are not only smarter, faster, and more efficient—but also fully private, equitable, and compliant.

Learn more about developing a security strategy for generative AI on AWS ›

Get started with Amazon Bedrock ›

Try out Amazon CodeWhisperer ›

Build and customize FMs on Amazon SageMaker ›

Elevate your security in the cloud with AWS ›