LTIMindtree

Guarding AI Innovation in the Cloud

Robust Framework for Responsible AI Adoption

Secure your AI models and data with proven AWS strategies

Table of contents

Executive summary

As generative AI becomes central to enterprise transformation, securing AI workloads is no longer optional—it's urgent. This whitepaper presents a comprehensive security framework for deploying AI on Amazon Web Services (AWS) cloud, tailored for enterprises navigating the complex intersection of innovation and risk. It is designed for CIOs, CISOs, cloud architects, and AI product leaders who are responsible for building secure, scalable, and compliant AI systems.

The paper begins by framing the growing imperative for AI security, highlighting that innovation without protection is unsustainable. The statistics in the whitepaper confirm that the need for proactive security is clear. Readers will discover how AWS's layered approach—spanning infrastructure, model, application, and data security—addresses these challenges head-on.

Real-world case studies, such as TaskRabbit's AI-enabled breach and Yum! Brands' ransomware incident underscores the financial and reputational risks of unsecured AI. The paper also explores emerging threats like shadow AI, prompt injection, and model poisoning, offering actionable strategies to mitigate them.
Regulatory frameworks, including the EU AI Act, NIST RMF, and ISO 42001, are unpacked to help organizations align with global compliance standards. AWS services like Amazon Bedrock and SageMaker are showcased for their built-in guardrails, encryption, and monitoring capabilities.

The best practices framework is provided to guide enterprises in applying zero-trust principles, securing model pipelines, and establishing AI-specific incident response protocols. The whitepaper concludes with a call to action: secure AI is not just a technical necessity—it's a business imperative. If your organization is adopting AI or scaling its cloud strategy, this whitepaper is your blueprint for doing so securely. Read the whitepaper to protect your innovation, earn stakeholder trust, and lead with confidence in the AI era.

Introduction
The growing AI security imperative

What if your most powerful AI innovation was also your greatest security liability? As enterprises race to embed generative AI (Gen AI) into their workflows, it is clear that success is not about who deploys AI first; it's about who does it securely. Any tech innovation without security isn't just risky; it's unsustainable. As AI workloads increasingly handle sensitive data, the real competitive edge lies in building trust, resilience, and compliance from the ground up. Data security is no longer a post-deployment consideration; it is a foundational requirement.

As Gen AI becomes central to digital transformation, the question is no longer whether enterprises should adopt AI but how to do so securely. Recent data paints a stark picture: between 2023 and 2024, 77% of businesses experienced data security breaches while using AI, and 70% of AI cloud workloads contained critical vulnerabilities.[1] These figures underscore the urgent need for a robust security framework that enables innovation without compromising enterprise integrity.

According to IBM's 2025 Cost of a Data Breach Report, 13% of breaches involved AI applications, with 97% attributed to inadequate access controls. The average cost per incident? A staggering USD 670,000. [2] [5] [9] Gartner further predicts that by 2025, 80% of unauthorized AI transactions will stem from internal access control failures and policy violations, not external attacks.3 This shift demands a new mindset—one that prioritizes holistic governance over perimeter-based defense.

Key security concerns in AI deployment

It is clear that security is the foundation of AI deployment, but with that foundation comes crucial challenges. As enterprises integrate AI into their workflows, four critical risks often surface:

1. Data exposure and model vulnerabilities

AI models require vast datasets for training and operation, and the continuous movement and processing of this data significantly expand the attack surface. Unintended access to sensitive resources, coupled with exploitable model behaviors, makes this one of the most pressing concerns in enterprise AI adoption.

01

Model poisoning and supply chain attacks

Adversarial actors can corrupt AI models by manipulating training data, creating backdoors or introducing biases. This risk is considerable when we use third-party models or data sources that require rigorous validation and monitoring throughout the product lifecycle.

02

Prompt injection attacks

Imagine an AI chatbot at a retail outlet mistakenly offering a high-value product for just $1. This scenario underscores the tangible risks of inadequate prompt security and the critical need for robust safeguards in AI deployment.

03

Data leakage and access control failures

In a leading reported incident, a major global enterprise accidentally entered sensitive semiconductor data into a public generative AI tool by feeding confidential source code and meeting notes into it. That incident led the organization to temporarily ban generative AI tools and then create in-house AI solutions to mitigate those risks. They also had to implement stricter internal protocols and restrict the data volume that can be fed into an AI application.

2. Infrastructure and deployment vulnerabilities

From development to production, the complexity of AI environments can expose various potential vulnerabilities, such as:

01

Misconfigured cloud resources

Research shows that over 3/4th of development teams use overprivileged default service accounts of Google's Vertex AI, and 91% of firms have risky admin access in their Amazon SageMaker notebook instances.[4]

02

Shadow AI proliferation

One emerging threat in the world of AI is "Shadow AI authorization." This threat involves employees' unsanctioned use of AI tools. Currently, only 37% of organizations have controls in place to detect and remediate shadow AI usage.[5]

03

Regulatory and compliance requirements

Like most new technology, the regulatory aspect of AI is expected to evolve and requires complex compliance requirements for organizations. The EU AI Act, the world's first comprehensive AI regulation, imposes strict requirements on high-risk AI systems with penalties up to €35 million.[6]

Key regulatory frameworks

1. EU AI Act

This legislation, applicable to all providers, categorizes AI-using apps by their risk level and then imposes equivalent security requirements on apps.

2. National Institute of Standards and Technology (NIST) AI Risk Management Framework (RMF)

Based on four key functions: Govern, Map, Measure, and Manage, this framework provides voluntary guidelines to identify, assess, and mitigate risks to ensure responsible AI deployment.

Figure 1: NIST AI risk management framework

3. ISO 42001

The globally accepted ISO 42001 is the first international standard for managing AI systems responsibly. This framework guides the establishment, implementation, maintenance, and improvement of an AI Management System (AIMS), ensuring ethically responsible development.

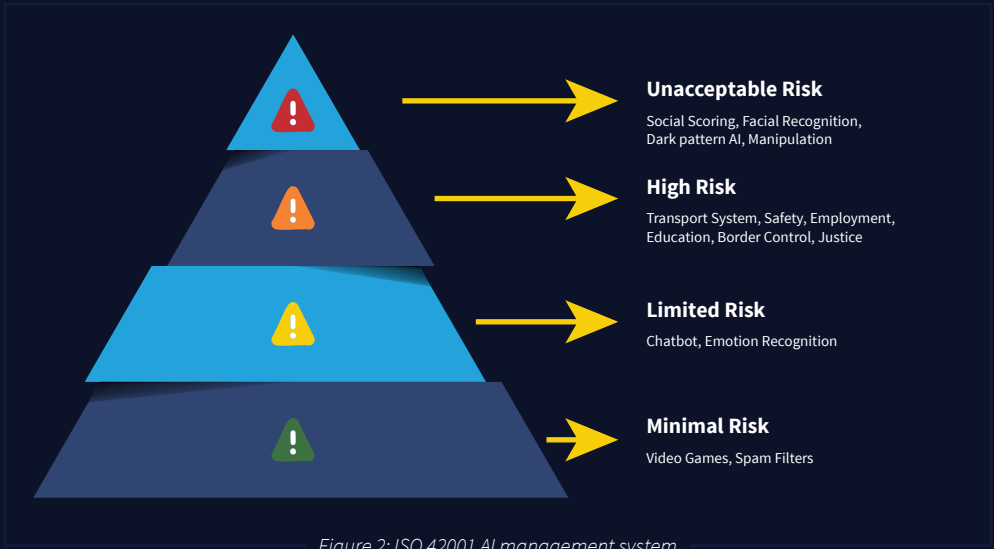

Figure 2: ISO 42001 AI management system

Compliance challenges

Maintaining compliance with technology that is still under development is never easy. Due to its unprecedented nature, regulators themselves are not sure of boundaries and often undercook or overcook compliance requirements.

Multi-jurisdictional complexity

For a similar workload, different regions have varying requirements. This multi-jurisdiction creates tough-to-deploy and complex region-wise matrices.

Rapid regulatory evolution

Due to the pace of marketing development and emerging threats, regulatory bodies' changes often outstrip organizational adaptation capabilities.

Technical implementation

Implementing hundreds of practices to ensure compliance by translating regulatory requirements into technical controls requires field experts, which many organizations lack.

Case studies on security incidents and financial impact

For many companies, AI is still in its early stages, but its risks have already been realized. From financial penalty to reputational damage, many companies have faced high-profile incidents.

1. TaskRabbit AI-enabled attack

In April 2018, TaskRabbit suffered a massive data breach that affected millions of personal and financial details. An AI botnet imposed a distributed denial-of-service (DDoS) attack and forced the company to shut down its operations for several weeks. This incident showcases the AI system's capability as a weapon of a business disruptor. [7] [8]

2. Yum! Brands ransomware with AI automation

In January 2023, Yum! Brands, a US-based fast-food corporation, faced a ransomware attack. This incident compromised the company's data, and their AI-powered system automated decisions to identify the target for data with maximum damage potential. This incident forced Yum! Brands close 300 UK branches for weeks. Such incidents demonstrate the impact of traditional attack vectors.[7] [8]

Financial impact analysis

The financial impact of an AI system breach goes beyond immediate response costs:

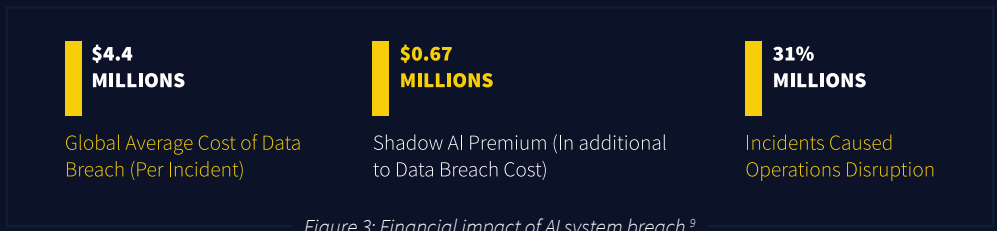

Figure 3: Financial impact of AI system breach [9]

AWS AI security framework
A comprehensive approach

AWS offers a suite of services that can address the AI deployment challenges and secure AI workloads throughout the lifecycle. Covering foundational principles, the AWS AI security framework provides robust protection without limiting innovation.

The Generative AI Security Scoping Matrix

This AWS application scoping matrix provides a framework that helps organizations assess and implement security controls throughout the Software Development Life Cycle (SDLC). This matrix categorizes workloads into five scopes:

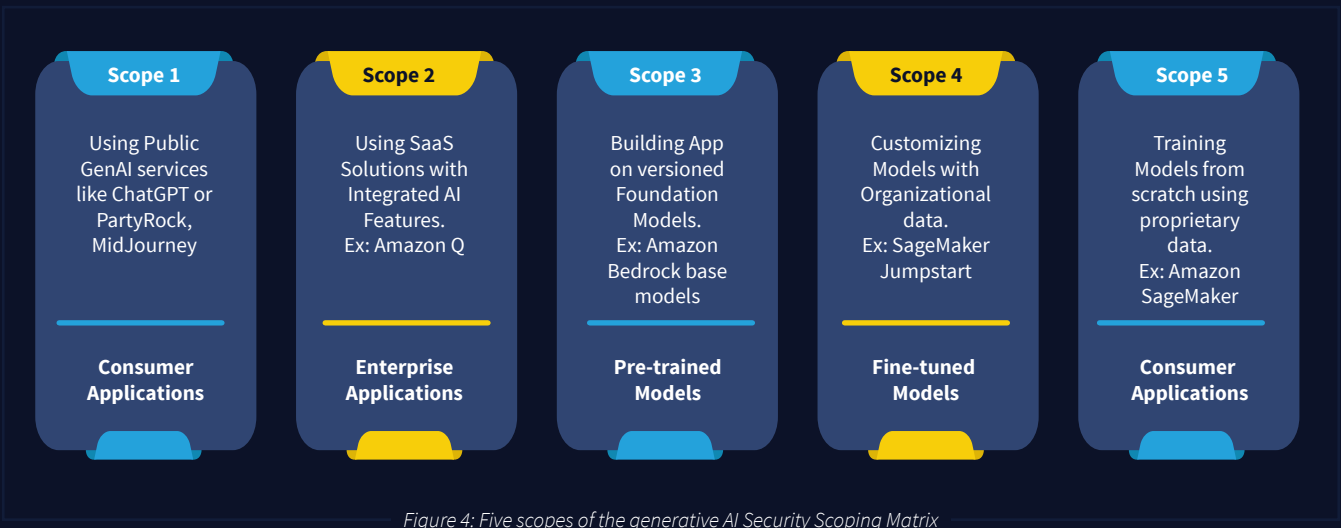

Figure 4: Five scopes of the generative AI Security Scoping Matrix

Depending on the application's nature, security considerations vary from basic governance for consumer apps to comprehensive model security for a self-trained system. This model should ideally be the starting point for all application builders.

Core security pillars

AWS secures AI infrastructure through four key pillars of AI security, each designed to address the unique challenges of AI deployment.

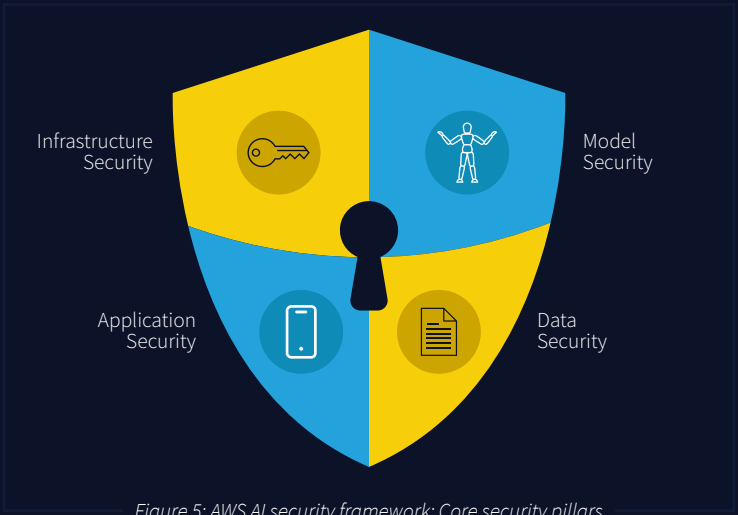

Figure 5: AWS AI security framework: Core security pillars

Infrastructure security

This pillar ensures that the underlying compute, storage, and network are robust, secure, and resilient. It leverages AWS's secure-by-design architecture and implements appropriate network isolation and segmentation

Data security

This crucial pillar of the AI security framework safeguards both the data that is consumed and generated by AI systems. The associated guidelines emphasize best practices for data categorization, classification, encryption, and access control implementation.

Model security

This pillar focuses on protecting AI models themselves by implementing access control methods, necessary encryption, and integrity monitoring for both pre-trained and custom models.

Application security

This pillar implements application usage security controls, including input validation, output filtering, and prompt injection protection.

AWS services and security features for AI workloads

1. Amazon Bedrock security configuration

As of 2025, Amazon Bedrock has a wide array of security features under its suite that can be configured to meet organizational custom security requirements:

Identity and access management

Using AWS IAM fine-grained access controls policies, usage models can be restricted based on user roles and the organization's policies. The following example demonstrates that the foundational model "Anthropic.Claude" is only accessible to the research team.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "bedrock:InvokeModel",
      "Resource": "arn:aws:bedrock:*:*:foundation-model/anthropic.claude-*",
      "Condition": {
        "StringEquals": {
          "aws:PrincipalTag/Department": "Research"
        }
      }
    }
  ]
}
```

Figure 6: Secured IAM policy for Anthropic.Claude users

Encryption and key management

By default, all data in Amazon Bedrock is encrypted in transit (TLS 1.2 minimum) and at rest (AES 256 encryption). Additionally, customers can opt for customer-managed keys for compliance and extra control over their custom models and fine-tuning data.

Network security

Amazon Virtual Private Cloud (VPC) endpoints can establish private connectivity between VPC and the Amazon Bedrock service without exposing data to the public internet. This practice ensures that API calls always remain within your trusted AWS network.

Model invocation logging

Enable comprehensive logging of all model interactions, including inputs, outputs, and metadata. This logging can be directed to Amazon CloudWatch Logs or Amazon S3 for analysis and compliance purposes.

2. Amazon Bedrock guardrails

Amazon Bedrock has the following guardrails that provide critical safety mechanisms to ensure responsible AI development:

Content filtering

This mechanism can filter harmful content such as hate speech, violence, sexual content, and misconduct by configuring thresholds. With varying sensitivity and strength levels, these filters can be applied to both prompts and responses.

Sensitive information detection

This feature can help identify and mask personally identifiable information (PII) and other sensitive data from responses. It supports both built-in detection parameters and custom regular expressions.

Topic filtering

This filtering prevents inappropriate discussion around specific subjects, persons, or things using natural language descriptions.

Prompt attack detection

By detecting attempts and proactively blocking them, this prevention system guards the AI workload against prompt injection attacks to manipulate system instructions through user inputs.

Contextual grounding

This measure reduces hallucination and ensures that model responses are relevant to user queries and are grounded in source material.

3. Amazon SageMaker security features

Amazon SageMaker has a broad suite of security controls for all AI models:

Network isolation

Private VPCs are a primary choice for deploying training jobs and interference endpoints to isolate networks. Restricted security groups and network access control lists (NACLs) should be used to control traffic flow.

Data encryption

AWS Key Management Service (KMS) keys can be used to encrypt all sensitive data at rest, such as training data, model artefacts, and interference data. Proper data-in-transit encryption should be enabled for data flow between SageMaker components.

Model registry and governance

SageMaker Model Registry can track model versions, approval workflows, and lineage. Automated governance should be implemented to ensure that only approved models are used for production.

Bias detection and explainability

For responsible AI practice, SageMaker Clarify detects bias in training data and model predictions. This feature provides transparency and fairness in AI decision-making processes.

3. Additional AWS security controls for AI workloads

Amazon GuardDuty for AI protection

A new feature of GuardDuty includes AI-specific threat detection capabilities. Amazon GuardDuty can now identify activities such as unusual API usage patterns and potential data exfiltration attempts.

AWS CloudTrail integration

CloudTrail is a one-stop solution for recording all API logs across all AWS services. This detailed audit trail can be helpful in compliance and forensic analysis. With CloudTrail Lake, advanced queries can be handled for large-scale AI audit requirements.

Amazon Macie for data discovery

It's advisable to use Amazon Macie to discover sensitive data and automatically classify it in your dataset. Macie ensures appropriate measures are in place before training AI models.

AWS Config for compliance monitoring

AWS Config rules can be used to ensure that best practices and compliance requirements are continuously being followed.

Best practices framework for secure AI deployment

1.

Apply zero-trust principles to AI workloads by

- Mandating authentication and authorizations at every step of the AI service intersection
- Implement network micro-segmentation around AI workloads
- Implement a system for continuous monitoring and validation of the AI system behavior
- Place appropriate validation in place to validate the sanity of AI outputs

2.

Establish comprehensive data governance by:

- Classifying data based on sensitivity and regulatory requirements
- Implementing data lineage tracking
- Implementing the proper encryption method for both data at rest and in transit
- Explicitly create a data retention and deletion policy for AI workloads

3.

Protect the security and integrity of AI models through

- Secure model development pipelines, model versioning, and approval workflows
- Monitoring model drift and degradation
- Protecting against adversarial attacks and model poisoning

4.

Develop AI-specific monitoring and incident response by

- Establishing a benchmark for normal AI operations
- Testing and updating AI security controls at regular intervals
- Creating an incident response playbook for AI security events
- Implementing real-time monitoring of AI system behavior

5.

Establish AI governance and compliance by:

- Creating responsible AI usage and ethics policies
- Implementing regular AI security assessments
- Maintaining compliance with the appropriate regulatory bodies
- Providing training to the development and operations team

Conclusion: Building a secure AI future

The adoption of secure AI practices requires a shift from the traditional approach. In the case of AI applications, unique security threats can emerge from both inside and outside the organization. Hence, it demands moving beyond regular perimeter-based security to a more comprehensive AI-aware security framework.

While AWS can provide services and features to protect foundational infrastructure, an organization's ultimate success depends on implementing the proper guardrails and security strategies for people, processes, and technology. The security practices outlined in the whitepaper are supported by LTIMindtree's extensive experience, AWS frameworks, and industry standards. They provide a roadmap to organizations looking forward to adopting AI, but are also skeptical about their security posture.

AI will continue to evolve and soon become integrated into various business operations. Organizations that would prioritize security now will be best positioned to realize the full value of AI while protecting their assets. The cost of implementing AI security is relatively low when compared to the potential impact of security incidents and the penalties associated with them. A compromised AI system not only shakes the confidence of stakeholders but also limits the speed of innovation. Proactive security of AI is not a technical necessity but a business imperative.

The future will not belong to those who keep their data secure and live with outdated technology, nor will it belong to those who adopt innovation without caring about data security. It would rather belong to those who can innovate fast but responsibly. Hence, embracing the power of AI while maintaining customers' trust, regulators' compliance, and stakeholders' confidence is essential. By implementing best security frameworks, organizations can confidently accelerate their AI transformation journey and sustain success.

Ready to innovate securely with AI?

Explore a proven framework that empowers enterprises to deploy generative AI with confidence. Learn how to mitigate risks, meet compliance, and build resilient AI systems—without compromising speed or scale. Write to us at *info@ltimindtree.com* to start your AI journey.

Author bio

Ashutosh Dixit
Principal Director of Cloud & Infra Consulting, LTIMindtree

Ashutosh, a passionate technology evangelist and AWS Ambassador, is a Principal Director at LTIMindtree. He leads the Cloud Strategy, Advisory, and Consulting team. He provides CIO advisory services worldwide and is an experienced solutions architect with expertise across hyper-scalers, FinOps, security, and adoption frameworks. As a renowned thought leader, he often speaks at industry events.

Citations

[1] *AI cloud workloads face greater critical security risks, Melvin Hipolito, securituybrief, July 01, 2025:*

https://securitybrief.com.au/story/ai-cloud-workloads-face-greater-critical-security-risks

[2] [5] [9] *IBM Report: 13% Of Organizations Reported Breaches Of AI Models Or Applications, 97% Of Which Reported Lacking Proper AI Access Controls, Michele Brancati, newsroom.ibm, July 30, 2025:*

https://newsroom.ibm.com/2025-07-30-ibm-report-13-of-organizations-reported-breaches-of-ai-models-or-applications,-97-of-which-reported-lacking-proper-ai-access-controls

[3] *2025 Gartner® Market Guide for AI Trust, Risk and Security Management, proofpoint:*

https://www.proofpoint.com/us/resources/analyst-reports/gartner-market-guide-ai-trism

[4] *AI is putting your cloud workloads at risk, George Fitzmaurice, itpro, March 19, 2025:*

https://www.itpro.com/cloud/cloud-security/ai-is-putting-your-cloud-workloads-at-risk

[6] *Article 99: Penalties, artificialintelligenceact.eu, August 02, 2025:*

https://artificialintelligenceact.eu/article/99/#:~:text=3.,financial%20year%2C%20whichever%20is%20higher

[7] [8] *Impact of artificial intelligence on criminal and illicit activities, dhsgov, 2024:*

https://www.dhs.gov/sites/default/files/2024-10/24_0927_ia_aep-impact-ai-on-criminal-and-illicit-activities.pdf