Registered in Germany Europa Revue Journal
ISSN: 1862-6564 Vol. 6 No. 9 (2026) Published September 30, 2026

CYBERSECURITY CHALLENGES IN THE ERA OF ARTIFICIAL INTELLIGENCE AND CLOUD TECHNOLOGIES

Samir Bakhtiyarov

Azerbaijan State University of Economics (UNEC)

BSc in Information Technology; MBA in Risk Management and Engineering

Abstract

Artificial intelligence and cloud computing have expanded the capabilities of modern information systems while simultaneously changing the attack surface that organizations must defend. Cloud platforms concentrate data, identities, APIs and computing resources in highly connected environments, while AI systems introduce models, training data, prompts and automated decisions as additional assets requiring protection.

Keywords

cybersecurity, artificial intelligence, cloud computing, zero trust, AI risk, data protection, incident response, digital resilience.

Introduction

The cybersecurity problem is therefore no longer limited to protecting servers and endpoints. Organizations must secure identities, cloud configurations, data pipelines, model interfaces and third-party services at the same time. The growing use of generative AI further complicates this environment because employees can transmit sensitive information to external services, attackers can manipulate model inputs, and automatically generated content can accelerate social-engineering attempts.

This thesis examines the principal cybersecurity risks created by the convergence of AI and cloud computing and proposes a risk-based control structure suitable for modern enterprises.

Main Part

How AI and cloud computing change cyber risk

Cloud adoption replaces a clearly bounded internal network with distributed services accessed through identities and APIs. Misconfigured storage, excessive privileges, exposed credentials and insecure interfaces can therefore become as

important as traditional malware. AI adds another layer because model behaviour depends on data, configuration and context, creating risks that conventional software testing may not fully capture.

At the same time, AI can strengthen defence. Machine-learning methods can help prioritize alerts, identify anomalous activity and support faster investigation. The important point is that AI is neither inherently a security solution nor inherently a security threat; its effect depends on how it is governed and integrated into the security architecture.

Risk area	Typical example	Priority control
Identity	Stolen cloud credentials	MFA, least privilege, conditional access
Data	Sensitive data exposed to AI service	Classification, DLP, encryption
Model/API	Prompt injection or abusive API use	Input controls, rate limits, monitoring
Cloud configuration	Public storage or excessive permissions	Configuration baselines and continuous scanning
Third party	Dependence on external model/cloud provider	Vendor due diligence and exit planning

Risk management principles

NIST's AI Risk Management Framework emphasizes lifecycle risk management and trustworthiness considerations. For enterprises, this means AI systems should be inventoried, assigned accountable owners, classified by impact and monitored after deployment. Cybersecurity teams should participate before production release rather than only after an incident.

A useful control model combines zero-trust principles with AI-specific governance. Every user, service and workload should be authenticated and authorized according to least privilege. Sensitive data should be protected independently of network location. AI applications should log important interactions, maintain model and version information, and define escalation procedures when outputs affect high-impact decisions.

Human and organizational factors

Technical controls are insufficient when employees do not understand how cloud and AI services handle information. Security awareness should therefore include practical rules on uploading confidential data, approving AI-generated code, verifying generated content and reporting suspicious outputs. Risk ownership must also be clear across IT, cybersecurity, business units, legal functions and external providers.

The strongest governance model is integrated rather than separate: cloud security, data governance, AI risk and incident response should share inventories, classification rules and escalation paths. This reduces gaps between teams and makes controls easier to audit.

Expanded Discussion: Threat Scenarios and Control Design

AI-cloud systems combine several trust boundaries. A user may access a cloud-hosted application that sends data to an external model, retrieves information from an enterprise database and returns generated content to another business process. Each transition creates a point at which authentication, authorization, validation or logging can fail. Security architecture should map these dependencies explicitly instead of treating the AI component as a self-contained application.

Identity remains central because cloud resources and AI services are frequently controlled through accounts, tokens and service identities. Strong authentication, least privilege, separation of duties and regular access review reduce the probability that a compromised identity can reach sensitive data or high-impact functions. Machine identities require the same discipline as human accounts because leaked API keys can provide direct access to automated services.

Data protection should cover both information supplied to AI systems and information produced by them. Enterprises should define which classes of data may be used with external services, how long prompts and outputs may be retained, and whether generated material can be reused in production processes. These controls become particularly important when employees use general-purpose AI tools outside centrally managed applications.



Figure 1. Cybersecurity control priority across AI-cloud risk areas.

Expanded Discussion: Monitoring, Incident Response and Proportional Governance

Continuous monitoring is necessary because the behaviour of an AI-enabled service can change even when the application code has not changed. Input patterns, underlying data, model versions and provider services may evolve. Security teams should therefore monitor unusual access, abnormal API consumption, data-transfer patterns and model-related incidents alongside conventional infrastructure events.

Incident-response plans should define who can disable an AI function, revoke credentials, isolate a data source and communicate with an external provider. The response process should also preserve logs and model-version information so that investigators can reconstruct what happened. Where AI outputs influence customers or important internal decisions, organizations should establish a path for correcting erroneous automated actions.

Governance intensity should be proportional to impact. A low-risk internal drafting tool does not require the same validation as a model that affects credit, employment, safety or significant financial decisions. Risk classification helps organizations allocate controls efficiently while avoiding the opposite errors of uncontrolled adoption and unnecessarily restrictive governance.

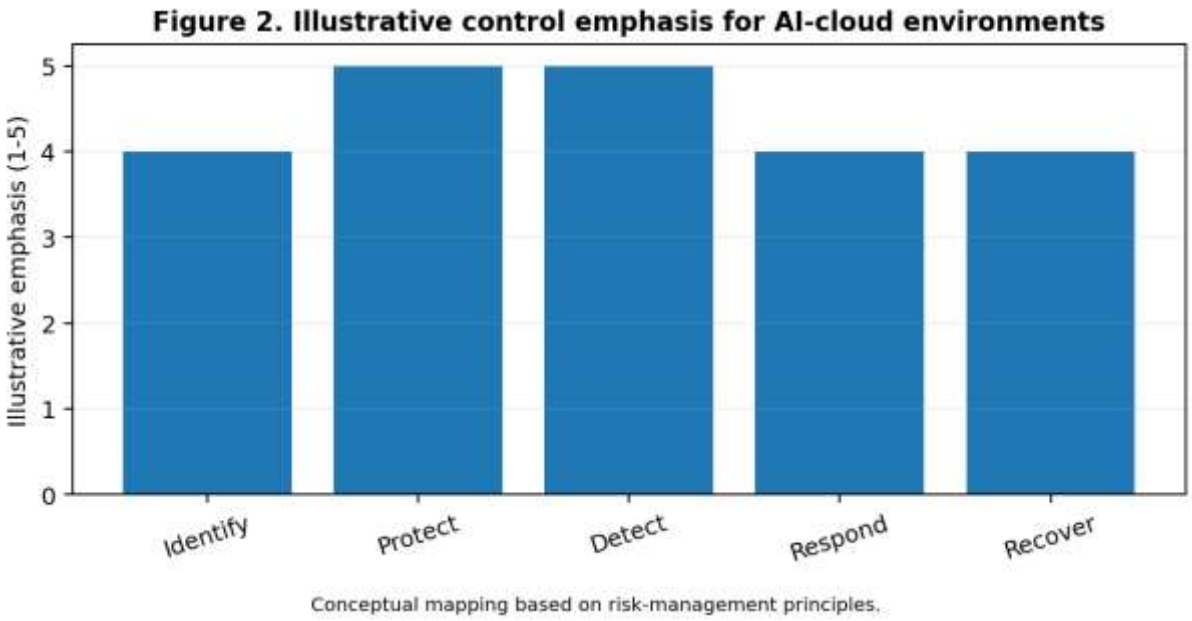

Figure 2. Illustrative control emphasis for AI-cloud environments.

Risk-Based Governance Model

A practical model can classify AI-cloud use cases according to data sensitivity, decision impact, external exposure and operational criticality. Higher-risk combinations should trigger stronger requirements for testing, approval, logging, human oversight and continuity planning. This approach connects technology governance with established enterprise risk-management practices and makes responsibilities clearer for both technical and business owners.

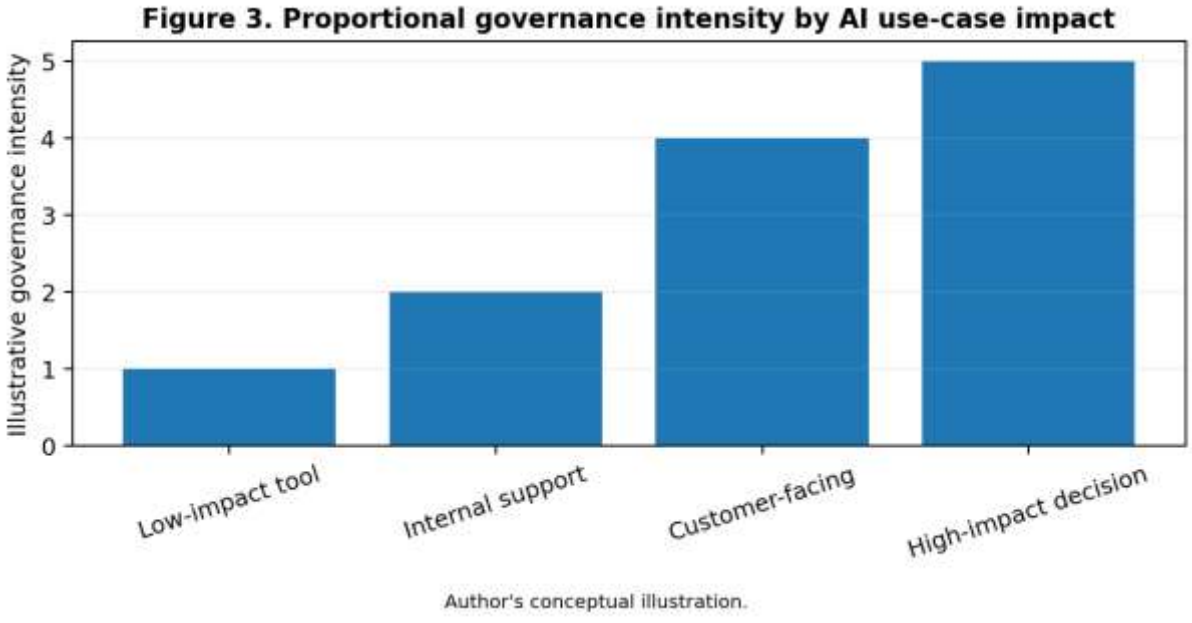

Figure 3. Proportional governance intensity by AI use-case impact.

Specific Enterprise Scenario: Securing an AI Assistant in a Multi-Cloud Environment

A practical example is an internal AI assistant that retrieves company procedures from a cloud knowledge base and uses an externally hosted language model to draft answers for employees. The main security boundaries are the employee identity, the application identity, the retrieval service, the document repository and the external model endpoint. A compromise at any of these points can expose sensitive information or permit unauthorized actions.

A zero-trust design requires explicit authentication and authorization for users and services rather than assuming that a request is trustworthy because it originates inside the corporate network. NIST SP 800-207 specifically shifts the security focus from network segments to users, assets and resources. For cloud-native and multi-cloud applications, NIST SP 800-207A further emphasizes granular application and service identities and policy enforcement through components such as API gateways and service proxies (Rose et al., 2020; Chandramouli & Butcher, 2023).

The AI-specific layer should add controls for prompt and output handling, sensitive-data filtering, model/version inventory, logging, human review for high-impact outputs and a defined method to disable the model connection during an incident. NIST's Generative AI Profile provides a useful basis for identifying risks unique to generative AI and selecting risk-management actions.

Framework Alignment

The proposed controls can be organized around the six functions of NIST Cybersecurity Framework 2.0: Govern, Identify, Protect, Detect, Respond and Recover. The addition of Govern in CSF 2.0 is especially relevant to AI-cloud systems because responsibility for model use, third-party providers, data handling and risk acceptance must be assigned before technical controls can be effective. This makes the cybersecurity model more specific than a generic list of safeguards: each control is tied to an enterprise risk-management function (Pascoe, Quinn, & Scarfone, 2024).

Conclusion

The convergence of AI and cloud technologies creates a cybersecurity environment in which identity, data, models and third-party services must be protected together. Organizations can reduce exposure by combining least-privilege access, secure configuration, data controls, continuous monitoring, AI lifecycle governance and employee awareness.

For managers, the key point is that rapid adoption should not precede risk architecture. Enterprises that establish governance and technical controls early can use AI and cloud services more confidently and scale successful applications without creating unmanaged security dependencies.

Adabiyotlar, References, Литературы:

1. Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. https://doi.org/10.6028/NIST.AI.100-1

2. National Institute of Standards and Technology. (2024). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1.

3. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero Trust Architecture. NIST SP 800-207. https://doi.org/10.6028/NIST.SP.800-207

4. Chandramouli, R., & Butcher, Z. (2023). A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Cloud Environments. NIST SP 800-207A. https://doi.org/10.6028/NIST.SP.800-207A

5. Pascoe, C., Quinn, S., & Scarfone, K. (2024). The NIST Cybersecurity Framework (CSF) 2.0. NIST CSWP 29. https://doi.org/10.6028/NIST.CSWP.29

6. Mell, P., & Grance, T. (2011). The NIST Definition of Cloud Computing. NIST SP 800-145. https://doi.org/10.6028/NIST.SP.800-145

7. OECD/BCG/INSEAD. (2025). The Adoption of Artificial Intelligence in Firms: New Evidence for Policymaking. OECD Publishing. https://doi.org/10.1787/f9ef33c3-en