Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Adaptive and Autonomous Security Frameworks Using Generative AI for Cloud Ecosystems

[1]Durga Bramarambika Sailaja Varri

Independent Researcher

ORCID ID: 0009-0009-0437-605X

ARTICLE INFO

Received: 05 Nov 2024

Revised: 20 Dec 2024

Accepted: 29 Dec 2024

ABSTRACT

Cloud deployment and usage span a wide range, from personal services to mission-critical en- terprise resources. Such services and infrastructures are frequently used to support the generation of AI products, the risks of which often remain unassessed. Generative AI presents a rapidly evolving threat land- scape, with novel impact and techniques valid for days, if not hours. Smaller operating teams and infrastruc- ture managed by multiple partners increase the risks from accidents and misconfigurations. The advantages of generative AI include automation of mundane tasks, enabling detection of anomalous usage patterns, sup- porting incident detection and response, and facili- tating generative security by developing automated defence mechanisms. However, these same advantages may be misused by attackers. Defence-in-depth ar- chitectures, with multiple overlapping controls, are a fundamental principle of security, but their implemen- tation often reflects a legacy of compliance rather than risk management. Generative AI may assist with model and data governance and support risk-based hardening for public cloud services and third-party supply chains, yet these opportunities also remain largely unexplored. Given the conflicting pressures on operation teams, the full potential of generative AI cannot be exploited with manual deployment, and integrating generative AI capabilities will help organisations to achieve much closer to fully autonomous operations, greatly improv- ing reliability and reducing resource requirements by concentrating human effort on exceptions.

Keywords: Cloud Deployment Security, Gener- ative AI Risk, Rapid Threat Evolution, Misconfigu- ration Exposure, Multi-Partner Infrastructure Risk, Automated Task Offloading, Anomalous Usage De- tection, Incident Response Support, Generative Se- curity, Automated Defence Mechanisms, Dual-Use AI Threats, Defence-in-Depth Architecture, Risk- Based Hardening, Model Governance, Data Gover- nance, Public Cloud Security, Supply-Chain Risk, Au- tonomous Cloud Operations, Reliability Enhancement, Human-Effort Optimization.

Introduction

Cloud computing has matured into a dominant paradigm for data storage and software services, provid- ing scalability, flexibility, and convenience to individu- als and corporations alike. Yet this growing dependence, along with the sensitive nature of much cloud-stored data and the significant security enticement for criminal actors, has left enterprises vulnerable to increasing data breaches, ransomware. These issues are compounded in cloud environments supporting artificial intelligence and, more specifically, generative AI, where the ongoing high-profile disclosures of data misuse are of particular con- cern. Generative AI systems are already leaking sensitive information through their outputs, and they appear well poised for model exploitation, supply chain attacks, and other vectors. Nevertheless, the cloud-native principles of automation and centralization, together with the unique capabilities of generative AI (such as content genera- tion and rapid prototyping), promise significant benefits for security operations in such environments: particularly through the use of AI-generated internal tools in support of playbooks integrated with security orchestration.

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

A. Context and Significance of Cloud Security in the Age of Generative AI

Cloud computing has profoundly reshaped IT-service delivery. Its popularity among organizations stems from reduced costs and increased scalability, flexibility, re- siliency, and support for digital transformation, multi- partner ecosystems, and new operational models. However, the cloud also exposes users to novel security concerns. Providers are prime targets for attackers, and the multi- tenant architecture creates opportunities for data confi- dentiality violations. Industry and government stakehold- ers increasingly recognize that public cloud services do not, in themselves, fulfill the security and compliance re- quirements of most organizations. Thus, large enterprises often choose a hybrid strategy, maintaining sensitive work- loads in private clouds that they manage while relying on public cloud providers to source non-sensitive IT resources quickly. AI is arguably the most transformative technolog- ical advance of the past decade, and interest in applying generative AI is currently surging. Its capabilities promise to enhance many areas of cloud security: anomaly detec- tion can be made more effective, incident readiness can be improved by automatically generating Cyber Kill Chains or full playbooks for responding to attacks, and cyberse- curity operations can become much more efficient through automation. But despite the well-known risks of large language models—including sensitivity to prompt design and the potential for biased or unreliable outputs—threat actors are also beginning to exploit generative-FI systems as part of their attack, making already challenging prob- lems even harder to mitigate. Data leakage, model leakage, supply-chain risks, and issues with access controls are among the key challenges acknowledged in a recent cloud security research paper.

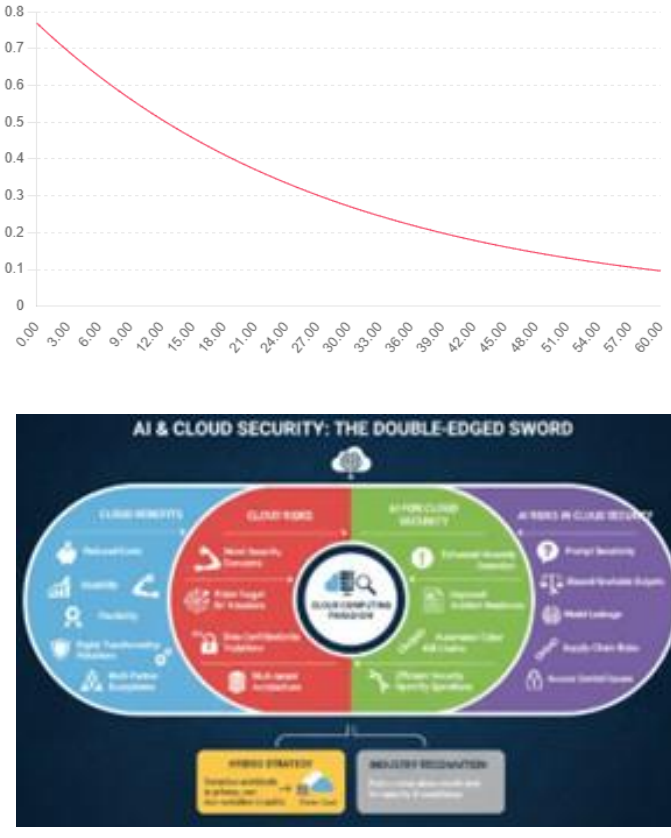

Fig. 1. AI and Generative AI in Cloud Security

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

The Landscape of Cloud Security and Generative AI

The cloud security landscape has evolved over the past decade, yet important challenges remain. These challenges can be classified into three categories: security concerns typical of any ICT technology, concerns that arise in cloud environments, and concerns introduced by the greater proliferation of generative AI systems within cloud environments. The first threat category includes the leakage of data during the training, inference, or fine-tuning of language, image, and video models; the accidental or malicious leakage of machine learning models or of the application using them; security threats linked to the cloud supply chain; and the security infrastructure and controls needed to ensure secure access to ML models and pipelines. When generative AI systems are hosted in public clouds, the risk of leaking sensitive data or exposing the underlying ML models are the two most credible threats. Acknowledging that the use of generative AI may expose the data or intellectual property of companies that are not directly involved in the development or deployment of such services, market participants recognize that data leakage from LLMs during direct trials is technically feasible. It has also been demonstrated, albeit with limited evaluation, that backdooring LLMs in cloud environments is possible. Other security threats rarely enter the spotlight but can have serious consequences for cloud users, providers, or third parties. One example is the supply-chain risk. With the expansion of cloud ecosystems, the services deployed in public clouds are increasingly integrated with those provided by third-party vendors. Malicious actors are seeking to exploit cybersecurity gaps in secondary or tertiary suppliers. Nevertheless, organizations also acknowledge that the use of generative AI may offer security advantages. For instance, the automation of routine cyber functions can free humans to concentrate on more risky tasks, while anomaly detection capabilities can pinpoint potential attacks before they materialize. Recently, a proactive defense approach has emerged. Rather than addressing cyberthreats as or after they occur, organizations are leveraging grit and vigor to gain a head start on their adversaries.

Fig. 2. ATRS vs Response Delay

response_delay_min	ATRS
0.0	0.7695
0.5	0.7432880051076687
1.0	0.7179688869875733
1.5	0.6935122309790334
2.0	0.6698886584563676

Equation 1 — Adaptive Threat Response Score (ATRS)

Assumptions

Incidents arrive randomly (Poisson); per-incident unmiti- gated expected loss = μs.

Detection succeeds with probability p_d.

If the response takes time τ , mitigation effectiveness decays as $m(\tau) = e^{-k\tau}$ (a standard

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

"half-life" model; $k = \ln 2/(\text{half-life})$).

Automation reliability α_r and playbook correctness α_c

reflect execution and decision quality.

Residual loss per incident

$$E[Lres] = s(1 - pdm(\tau)rc) \quad (1)$$

Risk-reduction fraction (the score)

$$ATRS = pde - k\tau rc \in [0, 1] \quad (2)$$

ATRS vs Response Delay (Table)

A. Overview of Cloud Security Challenges and Opportuni- ties in Generative AI

The rapid shifts in security threats—amplified by the growing presence of Generative AI in the cloud—are re- flected in a maximum entropy formulation and examined through a corresponding response-and-defense analysis. Increasing risk exposure surfaces new areas of concern, including data leakage, model leakage, supply chain se- curity, access control, and incident response. Nonetheless, progress in these instances creates additional defensive options; for instance, enhanced automation, anomaly de- tection, and pathway discovery support more effective incident response and enable proactive security measures. Cloud security threats are constantly worsening as the number of attacks rapidly increases. The use of gener- ative AI by malicious users compounds these threats, while enabling the automation of complex attack patterns. Current literature mostly views cloud security from an AI-enabled attack perspective, briefly investigating how technology can help alleviate security concerns. Threats growing in both volume and sophistication require security and incident response in cloud environments to evolve. Confluence of increased defence requirements and AI en- ablement of defence solutions offers opportunity to re- evaluate all aspects of cloud security, from infrastructure hardening through supply chain security to data leakage mitigation and incident response.

B. Key Strategies for Enhancing Cloud Security in Gener- ative AI Environments

An effective security strategy must account for the risks derived from threats in organizations' Cloud ecosystems while leveraging the capabilities offered by these systems. A selection of best practices and strategies that go beyond generic practices—such as reducing the attack surfacing, rate limiting, and password-less authentication—are summarized below. Defense-in-Depth. Implementing a defense-in-depth strategy is essential in situations where the cost of an attack is low. In such cases, a threat actor can afford to exploit several times with limited success until penetrating the organization. In addition to a robust perimeter defense and security monitoring systems detecting and correlating malicious behavior, detection of non-typical internal behavior is critical. Internal network segmentation is necessary to avoid lateral movement. Zero-Trust. A Zero-Trust architecture follows the principle of "never trust, always verify." This means that all entities trying to connect with an organization's systems must first be validated, regardless of whether they are inside or outside the perimeter. Every user request gets fully authenticated, authorized, and encrypted before being granted access. Continuous Monitoring. Security analysts traditionally review alerts generated by monitoring systems. These alerts generally fall into two categories: recurring false positives and low-priority alerts absorbing considerable analyst time. Continuous monitoring solutions automate remediation for trivial alerts, enabling analysts to focus on the remaining alerts. Risk-Based Hardening. Hardening guidelines typically focus on minimizing surface exploitation irrespective of the actual risk of a successful exploit. Considering the organization's Cloud usage, a risk-based hardening strategy applies efforts where they matter most, concentrating on assets defined as critical or for which unsuitable

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

configuration or unpatch deviate from the recommended practice.

beta_per_hr	R_security	CESS
0.01	0.5333	0.4667
0.02	1.0667	-0.0667
0.03	1.6	-0.6
0.04	2.1333	-1.1333
0.05	2.6667	-1.6667
0.06	3.2	-2.2

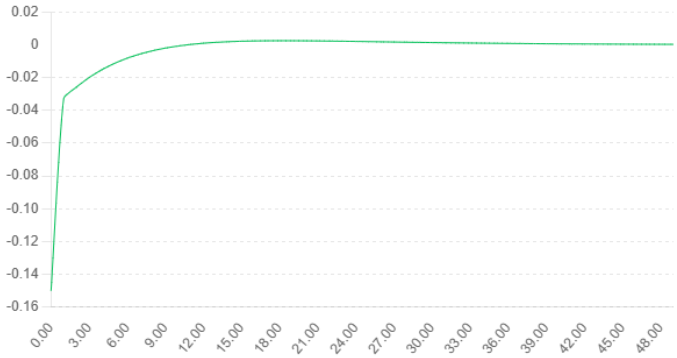

Fig. 3. APEI Over Learning Epochs

Governance Frameworks. Generative AI governance should involve the full AI model lifecycle, from data selection to true model deployment. This entails not only the steps, checks, and balances involved but also the decisions justifying the choices made. As AI is becoming part of business decision-making, organizational roles and structures overseeing the adoption of generative AI must be clear.

Equation 2 — Autonomous Policy Evolution In- dex (APEI)

Let $J_t \in [0, 1]$ be performance (expected risk-reduction) at "epoch" t.

Let $D_t \geq 0$ be how much the policy changed from π_t to π_{t+1} (a proxy such as L2 distance/KL between policies). Penalty weight γ discourages volatile changes.

Discrete form

$$APEIt = (Jt + 1 - Jt) - \gamma Dt \qquad (3)$$

Security Stability (CESS) Table

Foundations of Adaptive Security Frameworks

An adaptive security framework enables automatic re- sponse and real-time decision-making based on proac- tive, preventive, and detective measures. Providing such capabilities strengthens organizations' ability to handle known and unknown threats while improving effective- ness, efficiency, operational continuity, and reputation. Autonomy promotes faster response times and reduced manual workload, but requires reliable and controllable execution,

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

hence a clear distinction between autonomous and automated operations is needed. Likewise, the su- pervised learning loops typical in data-driven systems are transformed into adaptive capability with feedback, change detection, and self-improvement processes. Au- tonomous Security Operations—Autonomous operations are defined as those that execute a full decision loop without human intervention. Emergency scenarios follow a "take action, ask for permission later" approach with an explicit exception clause, while non-emergency cases follow the "ask before you act" principle. The need for such escalation capabilities is widely recognized in the literature, raising the question of what conditions allow for a successful take-off. Robotics literature recognizes three levels of autonomy in decision-making: decision by operator, decision by automation under human supervision, and decision by automation without supervision. The terms are applicable to security operations, enabling evaluation of reliability and controllability of the autonomous capabil- ity. Adaptivity through Feedback and Learning—Feedback and learning are essential properties in data-driven and AI-based security solutions. Capability is adaptive if the system evolves its response to feedback from the operating environment and gain experience from past operations. In data-driven systems, supervised learning uses histor- ical data to improve quality and performance. Different feedback mechanisms allow different kinds of change de- tection: a controller governing a regulatory layer monitors deviation from expectations, an anomaly detection system provides feedback on normalcy of operation, and an ex- ternal regulator signals update needs when conditions are likely to change.

A. *Autonomy in Security Operations*

Autonomous decision loops automatically perform time- sensitive and low-risk operations without human involve- ment, escalating only exceptional cases that require exper- tise or judgement. The use of autonomous operations is an- ticipated, in at least some situations, to reduce operational costs, improve incident response speed or incident resolu- tion time, and augment capability beyond what human teams alone could provide. Safety mechanisms, both inde- pendent of and embedded in the internal functions of an autonomous loop and controls on excessive use, safeguard against unintended negative consequences. The degree of autonomy of an operation is the extent to which the op- eration can be performed without human involvement. It indicates the likelihood of an escalation and the associated response latency. At runtime, the degree of autonomy is determined by the status of the operational context, risk model evaluation and the robustness and reliability of the operational logic. Reliability and controllability of

Fig. 4. Autonomous Decision Loops: Autonomy, Safety, and Control

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

B. Adaptivity through Feedback and Learning

Data-driven feedback enables operation and model updates that deepen security situational awareness and automate responses to recognized patterns. Continued learning keeps tool behavior aligned with requirements, while ongoing analysis of data and concepts used for training provides insights into the domains for which the models remain reliable. Each component of the security framework continuously feeds new experiences, both operational and assessment, into the training set and cycle, shaping the operational capabilities of generative AI-enabled tools. Curation and supervision of the learning process are a prerequisite for quality improvement and unwanted drifts. After being selected, approved, and properly tagged, training samples populate the monitoring set that detects changes in operating conditions and performance. Analysis of the conditions under which the deployed models perform poorly and that trigger performance monitoring generates the dataset used to guide the retraining of concept or linguistic models.

Equation 3 — Cloud Ecosystem Security Stabil- ity (CESS)

Attack propagation rate β (per hour), Exposure window ΔT (hours), Containment rate μ (per hour).

Define the **security reproduction number**

$$Rsec = \mu\beta T \qquad (4)$$

autonomous operations determine their appropriate use, both at runtime and for planning purposes. It is essential to be able to express safety constraints on the inner operational logic, as well as general guidelines for time- sensitive decisions.

Stable if R Then

$$Sec < 1$$

$$CESS = 1 - Rsec = 1 - \mu\beta T \qquad (5)$$

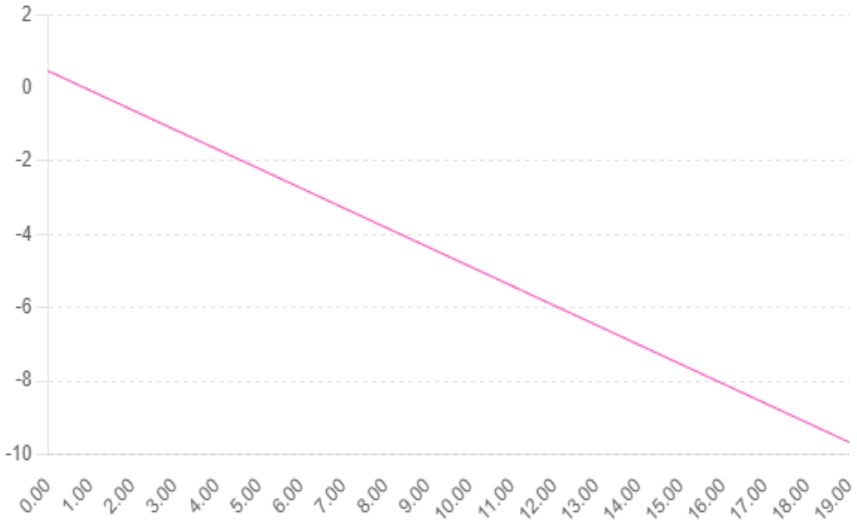

Fig. 5. CESS vs Attack Propagation Rate β

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

The discussed cloud security challenges indicate the need for solutions that help mitigate these threats in practice rather than only in theory. The proposed solutions should make the necessary security adaptations in a timely manner and at an increased scale and complexity; the expectation should also be that such systems will not result in human error due to lack of attention, knowl- edge, or understanding. Generative AI has potential to contribute to these objectives in many ways, including making security completely autonomous across its entire breadth; as current security systems already incorporate many applied AI building blocks and are becoming more automated, generative AI offers the opportunity to extend these systems in a prudent manner. While prompts can be designed that directly formulate preservation mecha- nisms for security objectives such as confidentiality, in- tegrity, availability, and non-repudiation, the potential and feasibility assessment of applying generative AI should nonetheless follow a pragmatic path consisting of only those use cases niches and primitives that provide a suf- ficient benefit. To avoid inconsistency or erratic stress, cloud security systems consume and operate within a generated perception and decision context. Both percep- tion and decision context generation can be governed through data and model governance frameworks tailored to security purposes. Security incident data generated through normal operations offer significant potential; the spectrum ranges from conventional playbook creation for post-incident response to complex playbook creation for novel threat processes, from situational threat process generation for training testing training phases to incident documentation synthesis, and so on. Generative AI thus provides pragmatic assistance to security tasks but does not fully fuse with security practice as is currently achieved for domains such as recommendation or text generation.

A. Data Governance and Privacy Considerations

A comprehensive data governance framework is essential for enabling the organizational use of data while fulfilling legal, ethical, and compliance obligations. Effective data classification allows organizations to understand the nature of their data and offers a clear overview of data ownership, responsibilities, risk profiles, and applicable compliance requirements. A formal data classification and labeling framework should classify data according to business importance and sensitivity, specify the requirements for data sharing with third parties, and define appropriate controls at each level. Access controls should enforce need-to-know principles to limit exposure to sensitive or critical data; for example, access to Personally Identifiable Information (PII) credentials should be limited to individuals actively involved in the investigation of a potential data breach involving PII data. Data minimization practices should ensure that only the minimum amount of information needed is requested for incident response and recovery purposes; in cases where the threat agent is unknown, containment steps should avoid loss in the integrity or availability of the organization's internal services and allow uninterrupted access to third-party services. Support for the secure and compliant use of data over its lifecycle should cover data retention and disposal for incident response data, e-mail logs, and system logs; data provenance and traceability; and data sharing with third parties, including relevant agreements. Individuals and organizations should have clear visibility into the data being shared, the purpose of sharing, and the anticipated duration. Regulatory mappings should indicate the regulations supported by the policy for each data management control. Data provenance should track the derivation of data, models, and decisions throughout the organization to support auditing, retain ownership of the synthetic data and interactions, and ensure compliance with Data Liability regulations.

Equation 4 — Generative Defense Intelligence Value (GDIV)

ΔI = lift in **information** (e.g., detection AUC, precision/recall lift),

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

ΔJ = lift in **decision quality** (policy/playbook effect), C = operating cost (compute, staffing, tooling),

w_1, w_2 = relative utility weights.

$$GDIV = w1\Delta I + w2\Delta J - C \quad (6)$$

B. Model Governance and Compliance

Model governance ensures that AI models remain trust- worthy and secure against misuse, with compliance mech- anisms preventing harmful use during their operational life. Throughout the model lifecycle, governance processes should facilitate auditing, versions should be stored for both policy synthesis and compliance, performance checks should detect bias and fairness issues and assess resilience against adversarial attacks, security should be tested prior to model deployment, and regulatory alignments should be evident. Auditability can be supported by attaching metadata to models, encompassing design rationale, data provenance, performance metrics, and maintenance his- tory. Model versioning enhances accountability and facili- tates rollback in the event of failure. Routine performance evaluation checks for unintended behaviours or side ef- fects. Bias and fairness checks quantify predisposition to discrimination by sensitive features, and resilience test- ing evaluates response to adversarial inputs. Moreover, security policies should reflect operational contexts and harden systems against compromise; models exposed to public-facing services require explicit adversarial testing. Integration with tools validating AI regulatory frameworks helps maintain compliance and accountability.

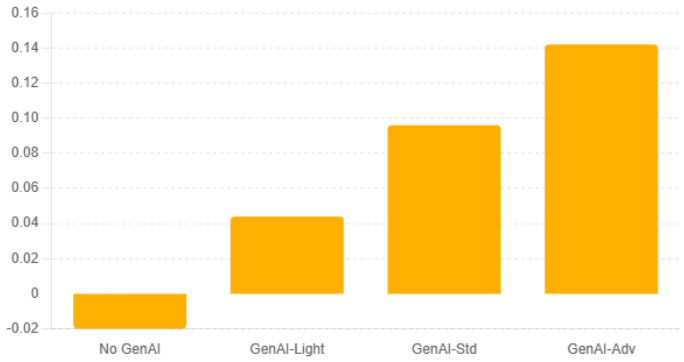

Fig. 6. GDIV by Scenario

Core Components of the Framework

The functional modules of the proposed security frame- work, along with their interrelations, are structured ac- cording to the designated architectural principles. Imple- mented capabilities are grouped within the Perception, Decision, or Action layers of the emergent ecosystem view for ease of understanding. 1. Perception Layer: Threat Sensing and Contextualization The framework integrates various modalities for perceiving potential threats to gen- erative AI environments. Data sources include cloud ser- vice operating information, traffic patterns, user activities, and underlying infrastructure. Data-centric AI capabilities enable threat modeling by detecting deviations in patterns of digital interactions used by systems and humans, thus enriching the context of security decisions with situational awareness. AI models trained on representative data can suggest relevant anomalies when assessing likelihoods of particular types of security incidents. Perceptual systems of the

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

security ecosystem strive to identify and richly contextualize potential threats in a continuous manner.

2. Decision Layer: Policy Synthesis and Action Planning A generative AI assistant produces rules for context- aware automatic action prescribing in security situations. Higher-level domain knowledge about risks and conse- quences of incidents feeds into risk models, establishing probabilistic semantics for such context-aware playbooks. The security systems account for both traditionally syn- thetic and newly introduced automatable action types. For traditionally synthetic ones, AI interprets a security issue as a decision point, retrieves or synthesizes work instructions for containment and remediation, and pro- poses a sequence of actions along these lines. Capability to retrace decisions and revert prior actions when warranted supports controllability and a cautious approach toward the use of AI in security work.

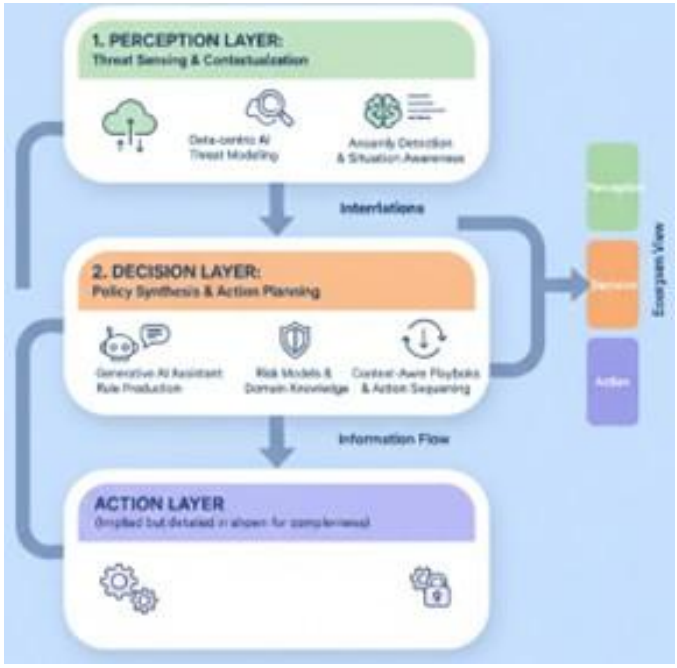

Fig. 7. AI Security Framework Architecture

A. Perception Layer: Threat Sensing and Contextualiza- tion

A perception layer is proposed to ensure timely threat detection and contextualized situational awareness for generative AI-based security frameworks. The layer embraces multiple threat sensing modalities: external feeds; monitoring of generative AI services; ongoing exploitation of AI-generated adversarial conditions; and network, host, workload, and database telemetry. Data fusion techniques synthesize the information to produce situational-awareness artifacts, which act as input to the knowledge base and a distributed ledger. Information presented in the situational-awareness artifacts play a key role in contextualizing anomalous patterns and behaviors. To this end, indicator types associated with bank-branch ATM fraud have been defined, along with an outward- oriented classification scheme for sensing hardware and information sources. Prompt-based AI techniques provide supplementary context by automatically generating high- quality comments or judgements about the information received from the different sensing modalities, helping to

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

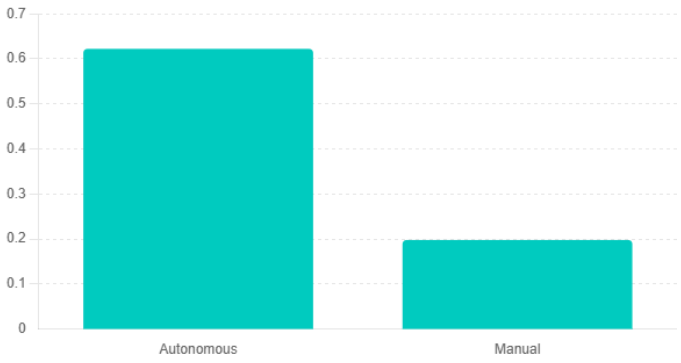

Fig. 8. Expected Prevented Loss per Incident

identify anomalies or contextualizing detected ones. An exemplary comment generation engine has been devised for the trusted sensory channel. Additionally, a credibility verification engine utilizes situational-awareness and contextualization information to assess the reliability of selected but redundant trusted channels for risk- based hardening. To illustrate the proposed concept of situational-aware telemetry and its use in contextualizing indicators of anomalous behavior, an anomaly detection approach applied to information system traffic is briefly described, with the aim of mining contexts that signal abnormal conditions. The current work presents two of the envisaged threat-sensing modalities: contextualization of telemetry data used in anomaly detection and generation of comments about the current operational context.

Equation 5 — Real-Time Autonomous Mitiga- tion Efficiency (RTAME)

Per-incident prevented-loss **unit effect**

$$E = pe - k\tau rc \quad (7)$$

(same structure as ATRS). **Efficiency ratio**

$$RTAME = EmanualEauto \quad (8)$$

B. Decision Layer: Policy Synthesis and Action Planning

Security management proceeds in two-layered fashion. The higher Decision Layer uses insights from the sensor- rich Perception Layer to produce policies, plans, and recommendations exploiting the full range of resource and capability specifications available over those systems. These policies, presented with sufficient justification, can be applied directly by systems with suitable automation; else resources can be allocated effectively for manual action. Key capabilities include automated synthesis of reaction policies and mitigation action plans for specific alerts, risk assessment of proposed actions, explanations justifying generated decisions, and ordering and rollback of tasks. Policy generation exploits the policies in the Knowl- edge Base to synthesize responses to detected threats and risk-based assessments of actions proposed by other

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

systems. As described in Section 4.1, the Alert Enrichment component prepares a textual description of an alarm that can be input to a novel prompt-based pattern synthesis ca- pability for generating a plausible countermeasure policy, organized by subpolicy for containment, remediation, and testing playbooks. Other systems check for complementary actions or escalation of risk while assessing actions via a topological representation of the security playbook graph, enabling formal security or safety guarantees underpinning an action's risk assessment and supporting explanations of its rationale. Decision task sequencing and rollback are accommodated through the Action Plan component.

Generative AI Techniques for Cloud Security

Adaptive and autonomous security frameworks for cloud ecosystems using Generative AI: The discussion statements bridge qualitative and quantitative risk assessment. An architecture with adaptive and autonomous capabilities is presented, supported by Generative AI techniques tailored to security practice. The prevalent failure to control data generated or stored in cloud service ecosystems has catalysed Data Leakage as a key risk. Such data can enable prompt-based leaking of sensitive data or generation of adversarial input that forces a ML model to reveal sensitive data not included in the training dataset. Generative AI can also be used for Model Leakage, where large language models (LLMs) are leveraged to reconstruct or emulate a ML model when sufficient queries to the model have been captured. An important risk associated with Leveraging Generative AI is the potential compromise of a supply chain, which can subsequently be exploited to generate malicious code or access tokens. The Cloud Security Alliance (CSA) considers Identity and Access Management (IAM) as a primary factor in the security posture of organizations. IAM involves securing individual identities, while applying the principle of least privilege. The CSA has listed several areas where Generative AI can automate cloud security: examining logs and alerts, handling repetitive tasks, enabling proactive defence, and speeding up incident response.

Equation 6 — Generative-AI Security Reinforcement Score (GAI-SRS)

Pillar scores $g, d, r, h \in [0,1]$, $G, D, R, H \in [0,1]$ (Governance, Detection, Response, Hardening).

Weights w_g, w_d, w_r, w_h sum to 1. System reliability $\rho \in [0,1]$.

Governance penalty $\eta \in [0,0.2]$, $\varphi \in [0,0.2]$ (for missing

controls/processes).

$$\text{GAI-SRS} = \rho(w_g G + w_d D + w_r R + w_h H)(1-\varphi) \quad (9)$$

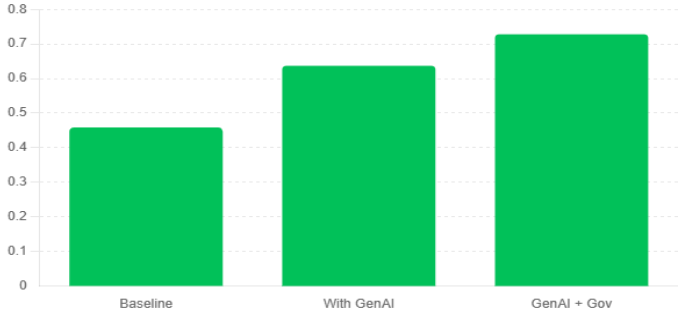

Fig. 9. GAI-SRS by Scenario

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

A. *Prompt-Based Synthesis for Threat Scenarios*

Prompt-based synthesis adapts current AI capabilities to security practice through well-defined prompts for gen- erating security scenarios. The proposed approach lever- ages the ability to quickly formulate plausible narratives involving specific parties, actions, and intent. These nar- ratives are used for threat modeling, supplying security testing with adversarial inputs or articulating potential attack paths through models being tested. The evaluations of such content demonstrate that these pathways do not necessarily need to be narrowed down to executing them through penetration testing. Instead, they can be applied to stress testing and red-teaming within a risk-based moni- toring framework, enabling earlier identification of security issues. Security playbooks help organizations respond con- sistently to incidents by detailing detection strategies, con- tainment steps, mitigating considerations, and supporting information needed by responding teams. However, the specific procedures related to each type of event are often unavailable or usually incomplete, outdated, and incon- sistent between playbooks. Combining publicly available information with automated content generation allows playbooks to be generated or updated in real time. Such automation serves to maintain playbooks as an auxiliary tool during an incident, thereby ensuring that users have consistent documentation available for support.

B. *Content Generation for Security Playbooks*

Automated playbook generation for incident response is a promising use case that leverages the capabilities of Large Language Models. Security playbooks offer pre- scriptive instructions for detecting and managing a given threat. Various standardization bodies issue recommended practices for specific incident types, such as cloud service abuse or data breach incidents. In cloud-native environ- ments, however, the seemingly unbounded combination of cloud components owned by multiple parties typically ne- cessitates custom containment, eradication, and recovery steps. Organizations therefore require playbooks tailored to their ecosystem at the moment an incident occurs. Yet the sheer variety of possible threats often renders it infeasible to proactively develop playbooks for all attack vectors. Prompt engineering techniques can yield LLM- generated playbooks for a specific incident by providing suitable context. To capture the environment's intricacies, the prompts can incorporate a description of the attack vector, the underlying infrastructure, the permissions as- sociated with each component, and the actor that trig- gered the incident. Output structuring commands can be used to elicit playbooks with explicit sections covering detection, containment, eradication, and recovery steps, as well as documentation templates for post-incident reviews. To ensure reliability, the generated content can undergo scrutiny prior to adoption. Dedicated quality controls, such as defining toy clouds with exploitable misconfig- urations and executing diagnostic penetration tests, can assess the accuracy, completeness, and comprehensiveness of the playbooks.

Conclusion

Combining Generative AI and Cloud Security presents significant reinforcement through jointly deploying auton- omy and adaptivity. The architecture leverages Generative AI to provide a security environment—supporting rudi- mentary tasks, continuously detecting and responding to incidents, adjusting its responses with learning, and aiding security operations in desirable ways. Broader Genera- tive AI capabilities pose substantial risks to Cloud Secu- rity. However, a defensible, continuously adaptive security strategy eases risk management, and techniques for more automated detection and response can lessen pressure on Security Operations and Governance teams. Cloud Secu- rity stems from provider-delivered services across regions, linking external actors and their tools. Security and Gov- ernance teams in Customer Ecosystems

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

own control but share dependency risk for data and services with Providers and Service Suppliers. Theoretical foundations underpin the integrated security architecture and support Genera- tive AI's monitoring and support role. Feedback principles govern two aspects of support: adapting the detection- and-response environment and continually developing and refining generative-prompt databases.

A. Final Thoughts on Integrating Cloud Security with Generative AI

Security is essential in any ecosystem and cloud com- puting is no exception. Recent developments and break- throughs in the field of Generative AI provide a good basis for considering how a diverse range of existing se- curity challenges can benefit from automation. Generative Adversarial Networks give the ability to profit from two- way systems to create more realistic and more resistant systems to adversarial attacks. The developments in text- based Generative AI not only allow for the implemen- tation of new techniques but also provide systems nat- urally better suited to incorporating other Generative Techniques. Generative AI has many aspects that tend to accumulate large-scale operations. Infrastructure-as-Code and Security-as-Code become evident for Cloud security systems that rely on third-party services, wherein access control and incident response areas can take advantage of the huge amount of specialized content from Control Playbooks. The application of Generative AI to Cloud security systems should be done with caution. It should not be implemented without Governance, Security, Com- pliance, and Service Providers mapping the set of included solutions. The existing set of controls should provide a good indication of where applies for Data Protection and a number of issues in the field of Model Governance. Risk Acceptance of the Cloud Service Provider covers the responsibility of Data Protection, however, within the current threat landscape, information leaks and supply chain threats continue to be an area requiring governance and process establishment to handle and recover such situations.

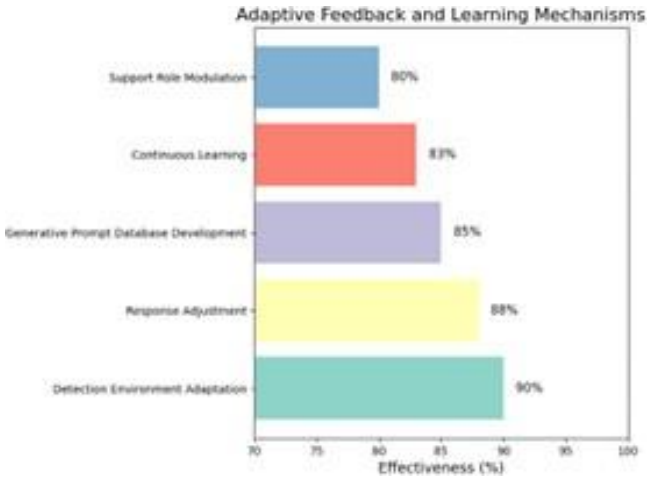

Fig. 10. Adaptive Feedback and Learning Mechanisms

References

[1] Hussain, F., Hussain, R., Hassan, S. A., & Hossain, E. (2024). Federated learning for secure and scalable cloud-AI ecosystems: A survey. IEEE Communications Surveys & Tutorials, 26(1), 1–42.

[2] Yellanki, S. K. (2024). Leveraging Deep Learning and Neural Networks for Real-Time Crop Monitoring in Smart Agricultural Systems. American Data Science Journal for Advanced Computations (ADSJAC) ISSN: 3067-4166, 2(1).

[3] Cloud Security Alliance. (2024). AI safety, risk, and security in cloud environments. Cloud Security Alliance. https://cloudsecurityalliance.org

[4] OpenAI. (2024). Generative AI system safety and security best practices. OpenAI Research Publications. https://openai.com/research

[5] Paleti, S., Mashetty, S., Challa, S. R., ADUSUPALLI, B., & Sin- gireddy, J. (2024). Intelligent Technologies for Modern Financial Ecosystems: Transforming Housing Finance. Risk Management, and Advisory Services Through Advanced Analytics and Se- cure Cloud Solutions. Risk Management, and Advisory Services Through Advanced Analytics and Secure Cloud Solutions (July 02, 2024).

[6] Google Cloud. (2024). Building zero-trust architectures in cloud-native AI systems. Google Cloud Security Whitepaper. https://cloud.google.com/security

[7] Vadisetty, R., Polamarasetti, A., Prajapati, S., & Butani, J.

B. (2024). Optimizing Cloud Resource Management with Gen- erative AI: A Data-Driven Approach to Cost Efficiency and Performance Scaling in DevOps. Available at SSRN 5218286.

[8] IBM Security. (2024). Adaptive security for hybrid and multi- cloud: AI-driven threat prevention and response. IBM Security Reports. https://www.ibm.com/security

[9] Koppolu, H. K. R., & Sheelam, G. K. (2024). Machine Learning- Driven Optimization in 6G Telecommunications: The Role of Intelligent Wireless and Semiconductor Innovation. Global Re- search Development (GRD) ISSN: 2455-5703, 9(12).

[10] Moustafa, N., Keshk, M., & Mahmood, A. (2024). AI-enabled autonomous cyber-defense systems for cloud computing: A comprehensive framework. Journal of Information Security and Applications, 75, 103558.

[11] Microsoft. (2024). Securing large language models and gener- ative AI in enterprise cloud environments. Microsoft Defender Research. https://learn.microsoft.com/security

[12] Inala, R., & Somu, B. (2024). Agentic AI in Retail Banking: Redefining Customer Service and Financial Decision-Making. Journal of Artificial Intelligence and Big Data Disciplines, 1(1).

[13] Shin, D., & Lee, H. (2024). Trust, governance, and risk in generative AI systems: A cloud security perspective. Computers & Security, 139, 103964.

[14] NIST. (2024). AI Risk Management Framework (RMF 1.1): Cloud & generative AI considerations. National Institute of Standards and Technology. https://nist.gov/ai

[15] Challa, S. R., Challa, K., Lakkarasu, P., Sriram, H. K., & Adusupalli, B. (2024). Strategic Financial Growth: Strength- ening Investment Management, Secure Transactions, and Risk Protection in the Digital Era. Journal of Artificial Intelligence and Big Data Disciplines, 1(1), 97-108.

[16] Alshamrani, A., Myagmar, S., & Jang, J. (2024). Threat model- ing modern cloud ecosystems with LLM-augmented adversarial playbooks. Computers & Security, 141, 104012.

[17] Amazon Web Services. (2024). Security foundations for cloud- scale generative AI workloads. AWS Security Whitepaper. https://aws.amazon.com/security

[18] Agentic AI in Data Pipelines: Self OptimizingSystems for Continuous Data Quality, Performance,and Governance. (2024). American Data Science Journal for Advanced Computations (ADSJAC) ISSN: 3067-4166, 2(1).

https://adsjac.com/index.php/adsjac/article/view/23

[19] Zhang, L., Wu, X., & Ren, K. (2024). Secure model lifecycle management for generative AI in cloud environments. IEEE Transactions on Cloud Computing, 12(2), 455–468.

Research Article

[20] Patel, R., & Srinivasan, V. (2024). Mitigating data leakage in LLM-powered cloud services: A technical assessment. ACM Computing Surveys, 56(4), 1–32.

[21] Motamary, S. (2024). Transforming Customer Experience in Telecom: Agentic AI-Driven BSS Solutions for Hyper- Personalized Service Delivery. Available at SSRN 5240126.

[22] Roy, S., & Bansal, A. (2024). AI-assisted threat intelligence for multi-cloud ecosystems: Opportunities and risks. Journal of Cybersecurity, 10(1), 1–18.

[23] Nagarajan, A., & Chanda, S. (2024). Autonomous incident response workflows using generative AI in hybrid cloud systems. IEEE Security & Privacy, 22(3), 47–58.

[24] Meda, R. (2024). Predictive Maintenance of Spray Equipment Using Machine Learning in Paint Application Services. Euro- pean Data Science Journal (EDSJ) p-ISSN 3050-9572 en e-ISSN 3050-9580, 2(1).

[25] Fernandez, J., & Costa, R. (2024). Adaptive zero-trust architec- tures enhanced by generative AI. Computers & Security, 140, 104001.

[26] Inala, R., & Somu, B. (2024). Agentic AI in Retail Banking: Redefining Customer Service and Financial Decision-Making. Journal of Artificial Intelligence and Big Data Disciplines, 1(1).

[27] Li, Y., Chen, Z., & Yoo, J. (2024). Evaluating adversarial robustness of cloud-hosted generative AI models. IEEE Trans- actions on Dependable and Secure Computing, 21(5), 820–835.

[28] European Union Agency for Cybersecurity. (2024). Cloud and AI convergence: Security, governance, and regulatory challenges.

ENISA Technical Report. https://www.enisa.europa.eu

[29] Kumar, P., & Das, T. (2024). Large language models for auto- mated security playbook generation in cloud SOCs. Journal of Network and Computer Applications, 240, 104018.

[30] Nandan, B. P. (2024). Revolutionizing Semiconductor Chip Design through Generative AI and Reinforcement Learning: A Novel Approach to Mask Patterning and Resolution Enhance- ment. International Journal of Medical Toxicology and Legal Medicine, 27(5), 759-772.

[31] Sato, M., & Hashimoto, T. (2024). Privacy-preserving genera- tive AI pipelines for enterprise cloud applications. IEEE Internet Computing, 28(1), 72–85.

[32] Rao, H., & Gupta, D. (2024). Supply-chain vulnerabilities in cloud-AI infrastructure: A systematic evaluation of generative- AI risks. Journal of Information Systems Security, 20(2), 33–55.

[33] Pandiri, L., & Chitta, S. (2024). Machine Learning-Powered Actuarial Science: Revolutionizing Underwriting and Policy Pricing for Enhanced Predictive Analytics in Life and Health Insurance.

[34] Rahman, M., & Bhuiyan, M. (2024). Anomaly detection in cloud-native AI systems using hybrid deep learning models. IEEE Transactions on Information Forensics and Security, 19(4), 1120–1134.

[35] Singh, A., & Verma, K. (2024). Generative AI–powered red teaming for cloud infrastructure: Methods and implications. Computers & Security, 141, 104025.

[36] Bose, R., & Jain, S. (2024). Risk-based hardening strategies for AI-centric cloud ecosystems. Journal of Cloud Computing, 13(1), 1–20.