Journal of Business and Management Studies
ISSN: 2709-0876
DOI: 10.32996/jbms
Journal Homepage: www.al-kindipublisher.com/index.php/jbms

JBMS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

A Lifecycle Governance Control Plane for Securing AI Workloads in Multi-Cloud Environments

Naresh Alapati[1], Ramachander Rao Thallada[2] and Koteswararao Nallabothu[3]
[1] *Principal Cloud Architect, Walmart, USA*
[2] *GRC Executive, Manulife, Canada*
[3] *Lead Software Engineer, Lowes, Charlotte, USA*
Corresponding Author: Ramachander Rao Thallada, **E-mail**: thalladaca@gmail.com

| ABSTRACT
Modern workloads supporting AI applications increasingly rely on multiple cloud platforms to enable growth, compliance, specialty services and resilience. This introduces a key challenge for governance in terms of dispersed security controls in the identity management system, model registry, deployment infrastructure, runtime endpoints, and monitoring stack. Cloud security techniques are traditionally provider-specific and infrastructure-centric, and thus fail to protect the end-to-end lifecycle of AI in multi-clouds. The present study proposes a multi-cloud AI governance framework, which includes the following three contributions: (i) MAGCP-6, a six-component governance control plane for AI applications in multi-clouds; (ii) LPEM, a lifecycle-based enforcement model for policies in AI in multi-clouds; and (iii) CGAL, a continuous governance assurance loop for AI in multi-clouds. These three components together constitute an integrated lifecycle-centric governance approach, which provides support for validating workload identity, making decisions based on policies, verifying model artifacts, attesting execution environments, monitoring runtime, and sustaining governance assurance.

| KEYWORDS

Governance; Control; Multi-Cloud; compliance; Artificial intelligence

| ARTICLE INFORMATION

ACCEPTED: 18 April 2026 **PUBLISHED:** 20 May 2026 **DOI:** 10.32996/jbms.2026.8.7.5

1. Introduction

Artificial intelligence (AI) workflows are no longer limited to a single cloud platform. Within a typical enterprise, training pipelines are run on a cloud platform or on-premises infrastructure, models are stored and promoted to a different platform, and inference endpoints emerge from a third cloud or hybrid platform. [1-3] The reason for this multi-cloud distribution is evident. Organizations adopt multi-cloud AI practices to gain access to unique hardware for computation, ensure jurisdictional data compliance, minimize vendor dependency, and increase resilience. However, multi-cloud flexibility creates a new category of AI governance problems that cannot be solved using cloud-native controls alone [4,5].

The root cause of AI governance problems lies in the fact that artificial intelligence is not governed at runtime alone. Compared to standard applications, AI pipelines include multiple stages that require trust assurance: data ingestion, preprocessing, model training, validation, artifact storage, deployment authorization, inference, monitoring, and decommissioning [6-8]. Governance gaps inevitably occur when multiple stages are scattered across different cloud environments. Identity authentication and enforcement are subject to changes depending on the cloud provider's policies. The integrity verification of models may be limited to registry boundaries. The deployment authorization may differ based on the location of the inference endpoints. Finally, the monitoring runtime behavior captures only one part of a larger picture [9,10].

An inherent contradiction exists between infrastructure security and AI governance. The former focuses on network security, role-based access, computing environments, and platform services, whereas the latter is concerned with protecting the

relationship of trust among datasets, models, identities, policy engines, execution environments, and runtime. When working in a multi-cloud environment, governance should become a portable control plane responsible for applying uniform and lifecycle-aware controls across all the involved clouds [11].

To resolve this problem, this study recommends adopting a governance-first approach to secure the distributed AI. This implies the integration of identity, policy, artifact verification, runtime protection, and observability into a cohesive model specifically designed to address multi-cloud AI governance. This involves three key innovations: The first is the MAGCP-6, a six-component multi-cloud AI governance control plane consisting of identity federation, policy enforcement, artifact integrity, execution trust, runtime threat monitoring, and assurance observability. The second contribution is a lifecycle-aware policy enforcement model called LPEM, which performs a joint assessment of identity trust, dataset classification, model provenance, execution environment trust, and runtime to make decisions. Finally, the third innovation is the continuous governance assurance loop (CGAL).

The objective of this study is to propose an architectural approach to govern the operation of AI models in a manner that allows them to operate safely in multi-cloud environments. This study is positioned at the intersection of cloud security, ML Ops governance, and trusted AI model supply chain.

2. Literature Survey

As recent advancements in AI governance have shown, one of the key directions is the transition from authorization techniques to runtime-aware policies. The example of runtime governance of agentic AI systems proves that the enforcement of policy requires consideration of the execution context, intermediate decisions, and environmental conditions, and not only predefined access rights and prompts [12]. Meanwhile, the development of continuous governance solutions suggests that telemetry-based trust monitoring frameworks can offer consistent visibility over deployed AI and provide verifiable assurances [13].

Moreover, recent security threats related to large language models make the implementation of runtime protection layers necessary. Recent research has revealed vulnerabilities in prompt injection, adversarial attacks, training-stage hacking, and leakage of sensitive information that cannot be resolved using infrastructure-level protection measures alone [14]. Moreover, threat analyses of generative AI reveal increased risks of automated phishing, malware generation, creation of adversarial content, and probing attacks in distributed AI, making runtime inference-time monitoring essential [15].

The complexity of governing multi-cloud environments includes heterogeneous identity management, inconsistent policy enforcement, and fragmented security-monitoring pipelines. Current research on hybrid-cloud security architectures reveals multiple issues related to the lack of cross-provider trust, regulation inconsistency, and vulnerability of inter-cloud communication, which prevent the unification of AI governance [16]. Meanwhile, cloud-native governance frameworks suggest that AI deployment requires the integration of policy enforcement, workload observability, infrastructure attestations, and jurisdictional authorizations in lifecycle-aware governance pipelines [17].

Data-centric governance research claims that secure AI deployment involves tracking dataset lineage, implementing adversarial machine learning protections, consistent red-teaming, and zero-trust enforcement in training and inference environments [18]. Additional security guideline frameworks expand upon this view by suggesting the structuring of governance controls regarding identity validation, artifact verification, assurance during deployment, inference time protections, and continuous monitoring compliance [19]. During inference, prompt injection detection solutions have revealed that runtime anomaly detection is essential for preventing manipulation and information theft from AI models [20].

Finally, AI governance in multi cloud environments can rely on the automated verification of compliance and anomaly awareness at the infrastructure level; however, this strategy usually does not incorporate provenance information and policy scores at different stages of the AI lifecycle [21]. In response, the framework of lifecycle-integrated AI–cloud security solutions recommend combining secure data pipelines, protected model supply chains, and runtime monitoring in the same cross-cloud trust architecture [22].

Supply chain integrity is a key requirement for securing AI deployment. Recent advances in blockchain-powered provenance verification frameworks have made tamper-resistant tracking of datasets, training lineage, and model promotions possible across distributed registries and deployment locations [23].

3. The Problem We Are Solving

The main concern addressed in this study is the lack of a unified governance framework applicable to AI workloads that are spread over different clouds. Currently, corporations use various governance strategies, including native IAM, registry-based verification, deployment policies and runtime monitoring. Although all these solutions work independently, they do not provide end-to-end governance for the entire lifecycle of an AI workload.

These solutions create various problems, which can be listed as follows: First, identity cannot be reliably trusted because a workload verified in one cloud might not be considered trustworthy enough for working with registries, deployments, and inferences from another cloud. The second issue is related to the lack of policy enforcement and reliance on phase-based policies: a workload may be compliant with local policies but not with external policies regarding the sensitivity of data, artifacts, or runtime behavior. In addition, there is an increased probability of degradation of model trustworthiness because of inconsistencies between clouds in terms of signature, provenance, and validation data. Another problem is associated with the detection of runtime threats, as only endpoint services are used to determine the presence of prompt injection, exfiltration, or probing without accounting for governance. Finally, the problem with observability and evidence collection is the difficulty in assessing whether the entire system is governed or only certain parts of it.

Thus, while the current solution exposes AI workloads to various threats, the real issue that needs to be solved relates to the non-portability and lifecycle awareness of the governance frameworks. The proposed architecture addresses this problem by providing a framework for the governance of AI workloads across clouds during their lifecycle, as explained in the following sub sections.

4. Implementation Architecture and End-to-End Governance Workflow

As illustrated in Figure 1, the proposed governance model is designed to be used as a portable governance control plane that operates on top of the AI distribution infrastructure without replacing any native services offered by the provider. Instead of utilizing a built-in cloud enforcement mechanism, the proposed model integrates the identity verification process, lifecycle-aware authorization, artifact validation, execution environment trustworthiness estimation, runtime monitoring, and telemetry-powered assurance process into a single governance flow across clouds. Therefore, the model can exist alongside cross-cloud ML Ops processes, promotions from model registries, inference engines running in containers, generative AI endpoints, CI/CD deployment pipelines, and hybrid training systems while maintaining underlying security measures. Architecturally, the suggested solution represents the combination of three different interconnected subsystems: MAGCP-6 (providing governance checkpoints based on the lifecycle), LPEM (context-based authorization based on lifecycle information), and CGAL (continuous governance assurance through monitoring).

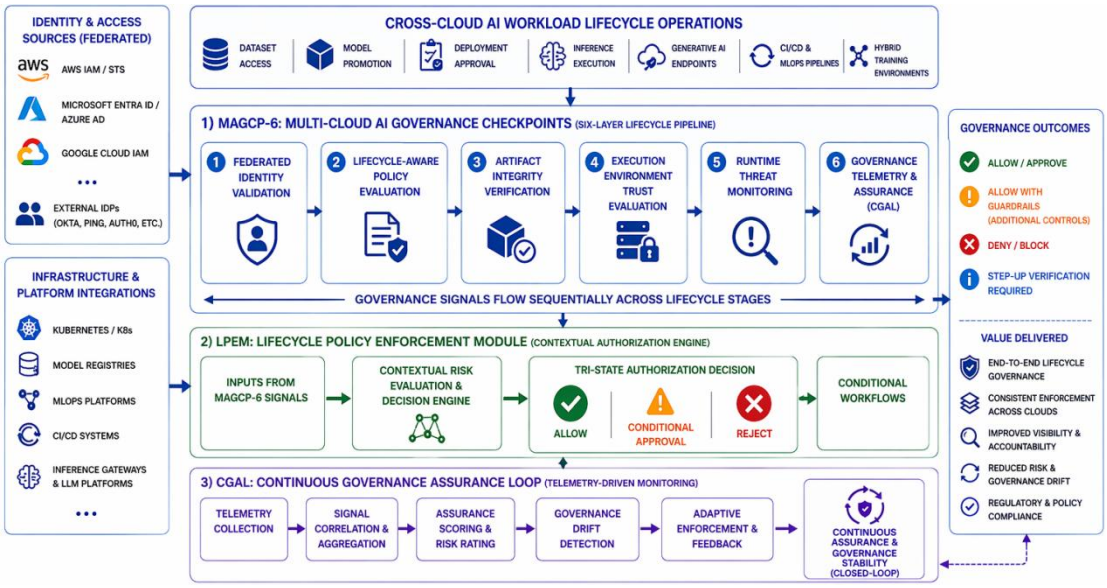

Figure 1. End-to-End Lifecycle Governance Architecture for Secure Multi-Cloud AI Workloads

4.1 Multi-Cloud AI Governance Checkpoints (MAGCP-6): MAGCP-6 constitutes the core of the governance model and results in the creation of six checkpoints that are laid out throughout the governance life cycle. The first layer of the enforcement mechanism checks the federated identity based on token validation, role translation consistency, fresh credentials, and workload provenance, considering all cloud providers, including AWS, Azure, and GCP. The second layer evaluates lifecycle-aware policies based on contextual signals in connection with classification levels for datasets, approval for artifacts, the security of the execution environment, and runtime anomalies. The third layer covers the aspect of artifact integrity and includes the verification of registry signatures, lineage validation, the relationship between datasets, and promotion eligibility. The fourth layer of the enforcement mechanism establishes trust in the execution environment through confidential compute attestation, secure storage of secrets, runtime isolation guarantees, and geographic locality for deployments. In the fifth layer, threats to runtime

are monitored by detecting malicious prompts, anomalous query frequencies, probing events, and data exfiltration signals. The sixth and final layer collects all telemetry from other checkpoints to maintain governance observability and enable drift detection using assurance score calculation.

4.2 Lifecycle-Aware Policy Enforcement Model (LPEM): Following identity validation, the Lifecycle-Aware Policy Enforcement Model (LPEM) acts as the decision-intelligence portion of the framework, whereby it processes information obtained at each stage to produce an authorization result that is grounded in contextual relevance. Rather than basing decisions on permission levels based on fixed roles, the model uses a lifecycle governance score based on weighted evaluations of identity validation, sensitivity analysis, provenance, execution environment attestations, and anomaly risk assessments associated with the governed process. Within its internal architecture, the LPEM converts each signal into a comparable trust score and then combines them based on a weighting scheme according to the level of sensitivity required for the workload and any legal or regulatory requirements. Finally, based on the overall governance score, the LPEM outputs one of three possible results: approval, conditional approval, or rejection of the proposed activity. In cases where execution receives conditional approval, the LPEM requires follow-up verifications before execution proceeds.

Before deployment, the framework checks whether the operating environment meets the set governance criteria based on workload sensitivity and infrastructure trust. This check entails an evaluation of the governance criteria for secret management, confidential computing attestations, runtime isolation, and geo-policy. Overall, these governance criteria are critical for ensuring that sensitive workloads are executed only in trusted infrastructures. Thus, deployment will consider not only artifact validation but also any risk factors in the execution environment.

After deployment, the governance of the runtime expands the governance criteria evaluation to cover the inference steps. Adversarial interactions at inference endpoints include prompt injections, suspicious query frequencies, model probing, and data leakage attempts. Any adverse interaction triggers an update of the governance criteria, where adverse interactions meeting certain thresholds result in measures such as access limitations, environmental validation, or runtime inference cessation.

4.3 Continuous Governance Assurance Loop (CGAL): The Continuous Governance Assurance Loop (CGAL) ensures consistent governance stability over time by transforming the concept of lifecycle enforcement into a feedback process based on telemetry, which includes the entire workload execution process. The CGAL uses telemetry data collected from identity federation services, artifact validation pipelines, execution attestations, runtime monitoring systems, and observability tools to derive the Governance Effectiveness Score, which is a measure of the alignment between the current trust posture and the governance policy posture defined earlier. In the case of any mismatches between the predicted governance posture and actual telemetry data, corrective actions are initiated, including, but not limited to, access restriction and policy modification. In contrast, within CGAL, telemetry data are collected, divergence from projected governance policies is identified, threshold values are compared with observed values, and remediation activities are triggered iteratively as a monitoring process to ensure the continuity of governance stability while evolving workloads operate within the distributed computing infrastructures.

The architecture can be deployed via a layered approach to integration that separates governance rules from provider infrastructure services but is compatible with enterprise AI frameworks. The federated identity broker allows integration with AWS IAM, Azure Entra ID, and Google Cloud IAM systems, whereas policy microservices enable lifecycle-aware authorization decisions using the LPEM framework. Artifact integrity adapters connect the logic of evaluating artifact integrity with the distribution of model repositories, and execution attestation agents allow integration with container orchestration systems such as Kubernetes. Runtime monitoring agents are integrated with inference gateways or service meshes, and telemetry aggregation services conduct CGAL-based drift detection and assurance scoring. Because all these integration points align with current DevSecOps and ML Ops practices, the adoption of lifecycle governance can be gradual and non-disruptive.

Among the most important implementation capabilities offered by the approach in question is the ability to propagate governance signals throughout the transition stages of the process, as opposed to assessing identity, artifact trust, deployment environment, and runtime execution in isolation. Identity trust affects the policy decision context, which in turn affects the eligibility for artifact promotion, which in turn affects the validation of the execution environment, and finally impacts the monitoring signals used during governance assessment. Such governance signals ensure that decisions made during the governance process remain consistent with the overall operational status of the workload.

5. Practical Deployment Scenarios in Enterprise Multi-Cloud AI Environments

AI systems on enterprise infrastructure are becoming more widely spread in terms of cloud deployments compared to the past, when they were limited to a single infrastructure boundary. Enterprises implement multi-cloud approaches to leverage GPU capabilities, adhere to specific regulatory requirements per jurisdiction, prevent vendor dependency, and provide better redundancy for crucial inference services. While this improves flexibility in terms of deployment strategy, it results in governance discontinuity in identity management, model cataloguing, deployment processes, and real-time monitoring capabilities.

From a practical perspective, this leads to inconsistent security controls across cloud providers, despite being consistent within each isolated environment. The framework introduced in this study provides a solution by providing a lifecycle-based control plane that enables trust continuity through the distribution of AI workflows. This section demonstrates how the proposed architecture can be applied to real-world enterprise environments in which AI processes are deployed across multiple heterogeneous infrastructures.

5.1 Scenario 1: Cross-Cloud Model Training and Deployment Pipelines

One common approach in enterprise deployments is to train models on one cloud and then deploy them to another cloud. For example, an enterprise might use a unique GPU offering from one cloud to train its models and then deploy the inference endpoints closer to its customers in another geographic location using another cloud service.

Model promotion pipelines are considered one of the important risks for governance in this approach because the source/identity of the artifacts may not be automatically propagated from the registry on the first cloud to that on the second cloud. The Artifact Integrity capability in MAGCP-6 ensures that the model signatures, lineage, and provenance are all checked before the model is promoted to the second cloud. With LPEM, this information becomes part of the authorization decision made during runtime, and any model without complete provenance cannot be further down the promotion pipeline.

Finally, the CGAL helps maintain governance consistency throughout runtime by monitoring telemetry signals through registry transitions and deployments.

5.2 Scenario 2: Federated Identity Across Distributed ML Ops Platforms

Enterprise ML Ops platforms frequently combine services from multiple providers. For example:

- Experiment tracking may run in one cloud
- Artifact storage may exist in another
- Deployment orchestration may operate in a hybrid Kubernetes environment

However, for such systems, the implementation of identity is often inconsistent because IAM rules native to each provider are not necessarily enforced across other cloud computing environments. To maintain consistency in trust, MAGCP-6 incorporates an Identity Federation module that implements cross-cloud validation for workload identity to achieve identity trust consistency across these services. Instead of depending on individual authentication barriers, this solution establishes a cross-provider identity map and assigns temporary credentials that can be applied throughout the entire process pipeline. As a result, LPEM adopts identity trust as part of its lifecycle-awareness decision model, where identity validation is not the only factor in deciding whether the system will approve the deployment.

5.3 Scenario 3: Secure Promotion of Models Across Registry Boundaries

Modern AI supply chains depend heavily on model registries that support version tracking, validation checkpoints and deployment approvals. However, when registries are distributed across cloud providers, promotion workflows become vulnerable to metadata inconsistencies and incomplete lineage verification.

The Artifact Integrity module addresses this issue by validating:

- Registry approval state
- Training lineage completeness
- Dataset provenance relationships
- Signature authenticity

before the models are promoted to production environments. Because these signals are fed directly into the LPEM decision engine, artifact trust becomes a mandatory requirement for lifecycle progression rather than an optional verification step. Consequently, model promotion pipelines remain governed even when registry services operate across heterogeneous infrastructure environments.

5.4 Scenario 4: Protecting Inference Endpoints in GenAI and Agentic Systems

Generative artificial intelligence systems pose additional runtime risks compared with traditional predictive systems. Prompt injections, probing of the machine learning model, and data extraction at runtime may occur post-deployment and cannot be prevented using static authentication controls. The Run-Time Threat Detection feature of the MAGCP-6 solution continuously analyzes the inference traffic for any anomalous patterns of interaction. The signals detected by this layer of detection were incorporated into the scoring performed on the CGAL telemetry.

When anomaly thresholds are exceeded, adaptive remediation actions may include:

- Restricting endpoint access
- Enforcing additional policy checks
- Requiring environment re-attestation
- Temporarily suspending inference sessions

This capability makes the framework particularly suitable for securing distributed GenAI deployments that span multiple cloud providers.

5.4 Scenario 5: Regulatory Compliance Across Jurisdiction-Specific Data Boundaries

Organizations operating in regulated industries must ensure that their sensitive datasets are confined to specific geographical locations and computing environments that are secure. This is made more difficult by using multi-cloud environments because data and models might shift across providers during transitions in the lifecycle of an application. With the Execution Trust component, geographical compliance checking becomes part of the deployment permissioning process. Before running the workloads, the system evaluates whether the selected environment conforms to the policies of the jurisdiction regarding the level of sensitivity of the datasets. Because all these indicators are assessed alongside provenance and identity assurance through the LPEM, there will be no issue in meeting the compliance requirements for deployment permissions.

6. Limitations, Design Tradeoffs, and Future Extensions

Architectures for governance that pertain to the management of multi-cloud AI systems must balance enforcement effectiveness and feasibility because lifecycle-aware security is applied throughout identity brokers, registries, runtime monitoring pipelines, and telemetry rather than isolated infrastructure-level controls. The proposed MAGCP-6, LPEM, and CGAL architectures facilitate trust assessment throughout the entire lifecycle but involve certain trade-offs pertaining to policy evaluation lag, identity federation, metadata standardization, runtime monitoring, scoring, and telemetry management.

Lifecycle-aware authorization allows for a more precise assessment because of dataset classification, artifact provenance, execution attestation, and runtime anomaly reporting, although this approach can increase the evaluation lag when processing high-volume transactions. Simultaneously, federating identities across AWS, Azure, and Google Cloud implies careful privilege management to prevent inconsistencies in cross-cloud authorization. In addition, metadata inconsistency is likely to pose difficulties in artifact provenance validation during the promotion of the model. The successful application of runtime monitoring relies on telemetry collection from inference endpoints. Moreover, weighting governance scores within the LPEM and CGAL requires appropriate threshold calibration according to workload sensitivity. Finally, the infrastructure costs associated with telemetry collection can be reduced through event-based assurance updates and sampling.

Agentic AI environments should be enabled by adding support for agent-level authorization, identity federation between agents, conversational anomaly detection, and multi-agent telemetry scoring.

7. Conclusion

As workload management in multi-cloud AI becomes more widespread, governance must extend beyond provider-specific areas. Identity enforcement, artifact validation, deployment authorization, execution protection, and auditing must function consistently throughout all stages of the lifecycle of a distributed AI system. However, many common frameworks focus on managing infrastructure and are fragmented over cloud boundaries, causing issues with trust continuity over the lifecycle.

In this study, a framework is introduced to govern multi cloud AI workloads via three interconnected layers: MAGCP-6, a six-layer architecture that ensures lifecycle governance through identity federation, policy evaluation, artifact integrity verification, execution attestation, runtime monitoring, and observability; LPEM, a lifecycle-aware authorization framework that incorporates dataset sensitivity, provenance verification, execution integrity, and anomaly detection in runtimes; and CGAL, a telemetry-based assurance loop that can detect governance drifts and adaptively address them.

These three layers create an abstraction for a consistent governance layer that is independent of the underlying infrastructure, yet is still compatible with any enterprise ML Ops or DevSecOps pipeline. In contrast to isolated checkpoint

enforcement, the propagation of governance-related signals across lifecycle boundaries enables trust consistency across the training pipeline distribution, registry promotion, multi cloud deployment, and inference execution.

This is especially relevant for enterprise settings, where regulation is needed, model development in the supply chain requires consistent life cycle governance, and generative AI models require reliable multi cloud AI infrastructure management tools. As distributed and autonomous AI systems gain momentum, lifecycle-aware governance, such as that proposed in this study, provides a solid starting point for their further development.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1]. Artificial Intelligence Risk Management, Framework: Generative Artificial, Intelligence Profile, NIST Trustworthy and Responsible AI NIST AI 600-1, https://doi.org/10.6028/NIST.AI.600-1
[2]. Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, https://doi.org/10.6028/NIST.AI.100-1
[3]. OWASP Top 10 for Large Language Model Applications, https://genai.owasp.org/
[4]. LLM01:2025 Prompt Injection, https://genai.owasp.org/llmrisk/llm01-prompt-injection/
[5]. A multilayer framework for good cybersecurity practices for AI, June 2023, ENISA, ISBN 978-92-9204-619-4, doi:10.2824/588830
[6]. AI Cybersecurity Challenges, December 2020, ISBN 978-92-9204-462-6, DOI 10.2824/238222
[7]. A COLLABORATION ACROSS INDUSTRY, ACADEMIA, AND GOVERNMENT, MITRE, ATLAS, 2023, https://atlas.mitre.org/pdf-files/MITRE_ATLAS_Fact_Sheet.pdf
[8]. J. Kressel et.al, SAFE-AI A Framework for Securing AI-Enabled Systems, April, 2025, https://atlas.mitre.org/pdf-files/SAFEAI_Full_Report.pdf
[9]. Visualizing Secure MLOps (MLSecOps): A Practical Guide for Building Robust AI/ML Pipeline Security, OpenSSF Whitepaper
[10]. Ashok Kumar Kanagala, Securing the AI Supply Chain: A Framework for AI Software Bills of Materials and Model Provenance Assurance, Transactions on Engineering and Computing Sciences; ISSN: 2756-2638, Vol. 14 No. 01 (2026) (119-129), https://doi.org/10.14738/tmlai.1401.19884
[11]. S M Zia Ur Rashid et.al, Lifecycle-Integrated Security for AI-Cloud Convergence in Cyber-Physical Infrastructure, arXiv:2602.23397v1 [cs.CR] 26 Feb 2026
[12]. Maurits Kaptein et.al, Runtime Governance for AI Agents: Policies on Paths, arXiv:2603.16586v1 [cs.AI] 17 Mar 2026
[13]. Eranga Bandara et.al, AI Trust OS — A Continuous Governance Framework for Autonomous AI Observability and Zero-Trust Compliance in Enterprise Environments, arXiv:2604.04749v1 [cs.AI] 06 Apr 2026
[14]. Pribisalić, M.; Martinčić-Ipšić, S. Security and Privacy of Large Language Models: Threat Taxonomy, Ethical Implications, and Governance. AI 2026, 7, 152. https://doi.org/10.3390/ai7050152
[15]. Yagmur Yigit, William J. Buchanan, Madjid G. Tehrani, Leandros Maglaras, Review of Generative AI methods in cybersecurity, Internet of Things and Cyber-Physical Systems, Volume 5, 2025, Pages 241-261, ISSN 2667-3452, https://doi.org/10.1016/j.iotcps.2026.04.001.
[16]. Sijjad Ali, Dhani Bux Talpur, Adeel Abro, Khulud Salem Alshudukhi, Ghadah Naif Alwakid, Mamoona Humayun, Farhan Bashir, Shuaib Ahmed Wadho, Asadullah Shah, Security and privacy in multi-cloud and hybrid cloud environments: Challenges, strategies, and future directions, Computers & Security, Volume 157, 2025, 104599, ISSN 0167-4048, https://doi.org/10.1016/j.cose.2025.104599.
[17]. Gagan Koneru,AI Governance Framework Adoption in Cloud-Native AI Systems: Phased Approach and Considerations, https://cloudsecurityalliance.org/blog/2026/01/27/ai-governance-framework-adoption-in-cloud-native-ai-systems-phased-approach-and-considerations
[18]. Didunoluwa Olukoya, Samson Onaopemipo Amoran, Oluwatosin Lawal, Malik Altawati, Saadat O Ibiyeye, Abdulaziz O Ibiyeye and Osondu C Onwuegbuchi. Data security and governance in the age of AI-enabled attacks. World Journal of Advanced Research and Reviews, 2025, 28(3), 1713-1722. Article DOI: https://doi.org/10.30574/wjarr.2025.28.3.4267
[19]. Ahmed Abugharbia et,al, Critical AI Security Guidelines, v1.2, SANS Research Program, https://cdn.lawreportgroup.com/acuris/files/Law-Report-Group-Files-New/SANS_Draft-Critical-AI-Security-Guidelines-v1_2.pdf
[20]. Lan, Qianlong & Kaul, Anuj & Jones, Shaun. (2025). Prompt Injection Detection in LLM Integrated Applications. International Journal of Network Dynamics and Intelligence. 4. 100013. 10.53941/ijndi.2025.100013.

[21]. Thapa, Ramesh & Kumar, Selva. (2025). Secure Multi-Cloud Architecture Using AI-Based Governance.
[22]. Valaboju, Praveen & Kadari, Rajeshwar. (2025). Integrating AI Governance Principles within Cloud Infrastructure: A Study of Apache CloudStack. Technix International Journal for Engineering Research. 12. 10.56975/tijer.v12i5.158051.
[23]. Adeyinka, Adetayo. (2025). Securing the AI Supply Chain: Using Blockchain For Verifiable AI Model Provenance on Government Clouds. SAMRIDDHI A Journal of Physical Sciences Engineering and Technology. 17. 2025. 10.18090/samriddhi.v17i01.05.