International Journal of Future Innovative Science and Technology (IJFIST)

|ISSN: 2454-194X | A Bimonthly, Peer-Reviewed, Scholarly Journal |

Volume 6, Issue 6, November-December 2023

DOI: 10.15662/IJFIST.2023.0606007

A Layered Security Framework for Enterprise Operations in the Generative AI and Agentic AI era in Regulated Cloud Environments

Lakshmi Kiran Meesala

Gilead Sciences Inc, NC, USA

ABSTRACT: The rapid enterprise adoption of generative AI (GenAI) and autonomous agentic AI systems has introduced a new class of operational risk that conventional security architectures were not designed to address. Unlike static applications, agentic systems plan, invoke tools, retain memory, and act across organizational boundaries with minimal human supervision, creating attack surfaces around prompt manipulation, tool misuse, excessive agency, and data exfiltration through model outputs. This paper proposes a layered security framework-spanning ingestion, processing, intelligence, and consumption tiers-purpose-built for organizations operating in highly regulated environments such as biopharmaceutical research, clinical operations, and manufacturing, where confidentiality, integrity, and auditability are non-negotiable. The framework integrates identity-centric access governance, real-time telemetry pipelines feeding modern SIEM platforms, model-aware guardrails, and policy-as-code enforcement points positioned at each agentic decision boundary. The design philosophy emphasizes least-privilege agent identities, deterministic tool gateways, continuous behavioral baselining, and compliance alignment with GxP, HIPAA, SOC 2, and ISO 27001 obligations. Implementation guidance covers deployment topology across hybrid AWS/Azure environments, telemetry routing through pipeline tools such as Cribl and DataBahn into SIEM platforms including Cortex XSIAM and Splunk, and automation of detection-rule lifecycle management. A case-study evaluation compares the proposed architecture against a traditional perimeter-and-SIEM baseline across detection latency, false-positive rate, containment time, and audit-readiness scores, drawing on three years of operational experience securing biopharma research, clinical trial, and manufacturing systems. Results indicate substantial improvements in mean time to detect and contain agentic-specific incidents, alongside measurable reductions in audit findings. The paper concludes with a discussion of scalability, residual limitations around emergent agent behavior, and directions for standardizing agentic AI security controls across regulated industries.

KEYWORDS: Generative AI security, agentic AI, SIEM architecture, zero trust, cloud security, regulated industries, threat detection, autonomous agents, identity governance, security operations

I. INTRODUCTION

Enterprises across regulated sectors are integrating generative and agentic AI into research, operations, and decision support at a pace that has outstripped the maturity of supporting security controls. In biopharmaceutical organizations specifically, AI systems now touch drug discovery pipelines, clinical trial data, manufacturing telemetry, and corporate knowledge bases-domains where a single integrity failure can compromise patient safety, regulatory standing, or years of intellectual property investment. The shift from passive GenAI assistants (which respond to prompts) to agentic AI systems (which plan multi-step tasks, invoke external tools, retain persistent memory, and act with delegated authority) fundamentally changes the threat model. An agent that can read a document, call an API, and write to a database is, from a security standpoint, a new kind of privileged identity-one that often lacks the audit trails, behavioral baselines, and access boundaries applied to human users or traditional service accounts.

This operational reality creates several converging challenges. First, data flows multiply: agents pull from research repositories, clinical systems, and manufacturing sensors, often crossing trust boundaries that were never designed for autonomous traversal. Second, traditional SIEM and detection content-built around known attack signatures and static rule sets-struggles to characterize agent behavior, which is probabilistic, context-dependent, and capable of legitimate-looking but inappropriate actions (e.g., an agent correctly authenticated but acting outside intended scope). Third, regulatory frameworks such as GxP, HIPAA, SOC 2, and ISO 27001 were not written with autonomous decision-

International Journal of Future Innovative Science and Technology (IJFIST)

|ISSN: 2454-194X | A Bimonthly, Peer-Reviewed, Scholarly Journal |

Volume 6, Issue 6, November-December 2023

DOI: 10.15662/IJFIST.2023.0606007

making systems in mind, leaving compliance teams to interpret existing controls-access governance, change management, audit logging-onto a class of system that did not exist when those controls were drafted.

These pressures motivate the need for a security architecture that treats agentic AI not as an application feature to be bolted onto existing infrastructure, but as a first-class identity and workload category requiring dedicated ingestion, monitoring, and governance layers. Drawing on three years of hands-on architecture and operations work securing cloud and on-premises environments in a regulated biopharmaceutical enterprise-spanning SIEM design and optimization, cloud security architecture across AWS and Azure, and the build-out of detection content and incident response processes-this paper makes the following contributions: (1) a layered reference architecture for securing GenAI and agentic AI systems across ingestion, processing, intelligence, and consumption tiers; (2) a deployment-level design covering hybrid cloud topology, telemetry pipeline routing, and SIEM integration suited to regulated environments; (3) an evaluation methodology and case-study comparison against a traditional security baseline; and (4) a discussion of practical trade-offs and open problems in scaling agentic AI governance across distributed enterprise systems.

II. ARCHITECTURAL FOUNDATIONS AND PROBLEM DEFINITION

2.1 Problem Formalization

Securing GenAI and agentic AI systems in an enterprise context requires solving a constrained optimization problem: maximize the operational value of autonomous and generative capabilities while bounding the probability and impact of unauthorized data access, unintended action execution, and compliance drift, under constraints of acceptable latency, infrastructure cost, and existing organizational tooling. Three classes of risk define the boundary conditions of this problem.

Input-layer risk arises from prompt injection, where adversarial or unintended instructions embedded in documents, emails, web content, or tool outputs alter an agent's planned behavior. Because agents frequently ingest unstructured content from multiple internal and external sources, the input layer cannot be fully sanitized using traditional signature-based filtering alone.

Decision-layer risk arises from excessive agency: an agent granted broad tool access to accomplish a task may use that access in ways the task did not require, particularly when planning is non-deterministic. This is compounded in multi-agent or tool-chaining architectures, where one agent's output becomes another agent's input, propagating errors or malicious instructions across a chain with no single point of human review.

Output-layer risk arises from data exfiltration or leakage through model responses, logs, or downstream artifacts, especially when models are fine-tuned or retrieval-augmented over sensitive corpora such as clinical trial records or proprietary research data.

2.2 System Constraints

The proposed architecture is bounded by constraints typical of regulated biopharma and similarly regulated enterprises:

- **Regulatory constraint**: All data handling, access, and change activity must be auditable to GxP, HIPAA, SOC 2, and ISO 27001 standards, requiring immutable logging and defined retention.
- **Latency constraint**: Security controls inserted at agentic decision points must not introduce unacceptable delay into research or operational workflows, typically bounding inline policy checks to low-double-digit millisecond overhead.
- **Heterogeneity constraint**: The environment spans hybrid AWS and Azure infrastructure, on-premises storage and backup platforms (NetApp, Veeam, Commvault, EMC), and a SIEM/telemetry stack (Cortex XSIAM, Splunk, Cribl, DataBahn), requiring the architecture to be platform-agnostic at the control layer.
- **Change-velocity constraint**: GenAI and agentic frameworks evolve rapidly; the architecture must support iterative onboarding of new models and agent frameworks without requiring redesign of the underlying security fabric.

2.3 Core Design Principles

The architecture is guided by four principles. Identity-centric governance treats every agent, tool, and model endpoint as a distinct identity subject to least-privilege access policy, never inheriting blanket human or service-account permissions. Deterministic mediation ensures that all agent-to-tool and agent-to-data interactions pass through policy-enforced gateways rather than direct, unmediated calls, preserving a consistent control point regardless of which model

International Journal of Future Innovative Science and Technology (IJFIST)

|ISSN: 2454-194X | A Bimonthly, Peer-Reviewed, Scholarly Journal |

Volume 6, Issue 6, November-December 2023

DOI: 10.15662/IJFIST.2023.0606007

or framework is in use. Continuous behavioral baselining replaces static rule-matching with profiles of expected agent behavior, flagging deviation rather than relying solely on known-bad signatures. Compliance-by-construction embeds audit logging, retention, and policy mapping directly into the data and control plane rather than retrofitting compliance evidence after deployment.

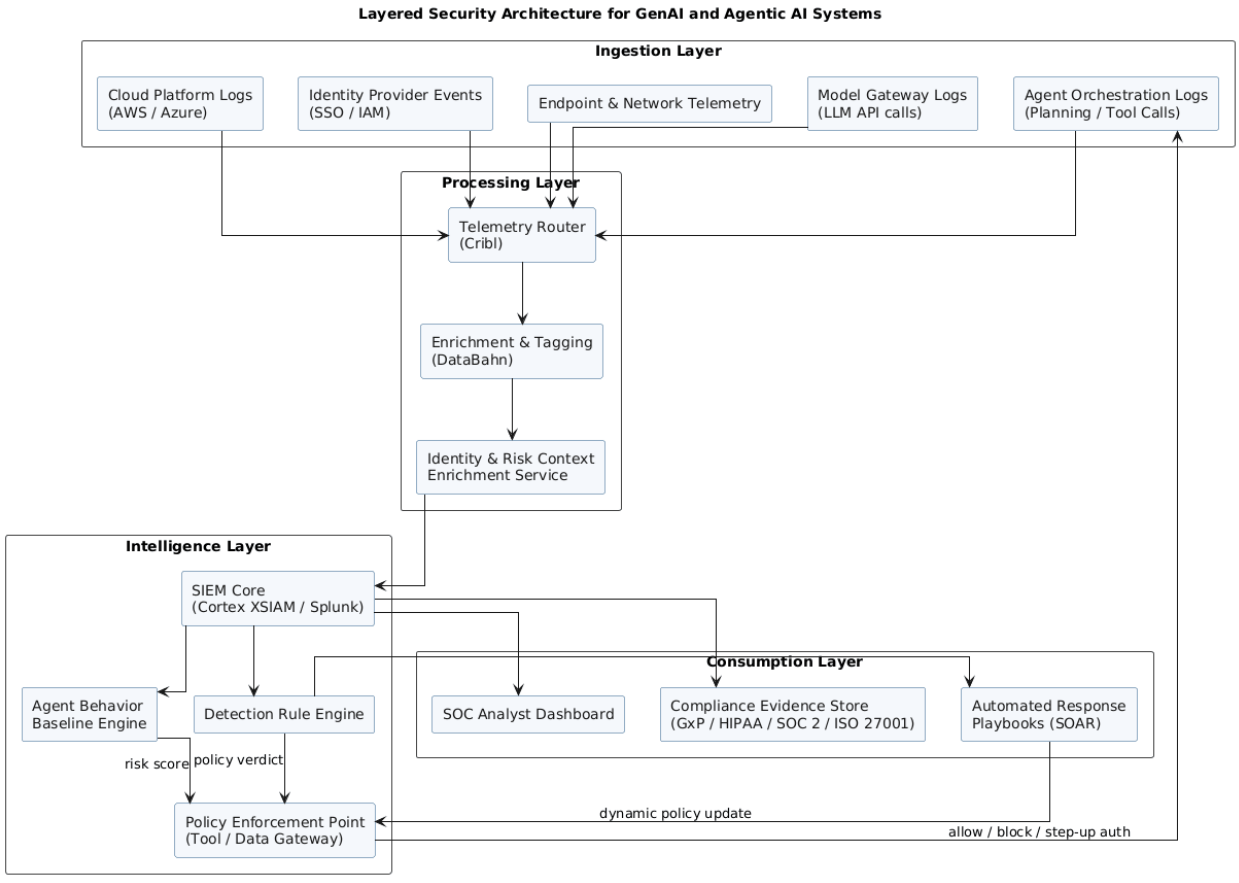
III. PROPOSED SYSTEM DESIGN

The proposed architecture is organized into four layers: ingestion, processing, intelligence, and consumption. Each layer addresses a distinct stage of the data and decision lifecycle, from raw telemetry collection through to the point where a human or downstream system consumes an AI-derived output or action.

3.1 Layer Overview

The ingestion layer collects telemetry from cloud platforms, identity providers, endpoint agents, model gateways, and agent orchestration frameworks, normalizing heterogeneous log formats before routing. The processing layer applies enrichment, filtering, and routing logic-implemented in this design using pipeline tooling such as Cribl and DataBahn-to reduce noise and tag events with identity and risk context before they reach storage or detection systems. The intelligence layer comprises the SIEM and detection content (Cortex XSIAM, Splunk) along with model-aware analytics that score agent behavior against baselines, in addition to the policy enforcement points that gate agent-to-tool and agent-to-data calls in real time. The consumption layer delivers detection output, compliance evidence, and dashboards to SOC analysts, compliance teams, and automated response playbooks.

3.2 Architectural Diagram



3.3 Subsystem Interactions

|ISSN: 2454-194X | A Bimonthly, Peer-Reviewed, Scholarly Journal |
Volume 6, Issue 6, November-December 2023
DOI: 10.15662/IJFIST.2023.0606007

The telemetry router, implemented via Cribl, performs initial volume reduction and format normalization across the five ingestion sources, allowing the downstream enrichment service to operate on a consistent schema regardless of whether an event originated from a cloud control-plane log, an identity provider, or an agent orchestration framework's tool-call log. This decoupling is essential in heterogeneous environments where AWS CloudTrail, Azure Activity Logs, and agent framework logs (e.g., from LangChain-, AutoGen-, or custom orchestrator-based agents) emit structurally incompatible event formats.

The enrichment and risk-context layer, implemented via DataBahn in this design, attaches identity metadata, data-sensitivity classification, and prior-behavior context to each event before it reaches the SIEM core. This step is what allows the downstream baseline engine to evaluate "is this agent action consistent with its historical and declared scope" rather than relying purely on static correlation rules.

The agent behavior baseline engine is the architectural component most distinct from a conventional SIEM deployment. Rather than matching events to known-bad signatures, it builds rolling behavioral profiles per agent identity-covering tool-call frequency, data-access scope, and inter-agent call patterns-and emits a risk score consumed by the policy enforcement point. The policy enforcement point (PEP) sits in-line between agents and the tools or data stores they call, executing allow, block, or step-up-authentication decisions in real time based on combined input from the baseline engine and the detection rule engine. This inline placement is what allows the architecture to intervene *before* an inappropriate action completes, rather than detecting it only after the fact in the SIEM.

Finally, the consumption layer separates SOC-facing detection workflows from compliance-facing evidence collection, recognizing that auditors and compliance teams require structured, retained evidence of control operation (e.g., proof that a PEP blocked an out-of-scope data access attempt) distinct from the real-time triage view used by SOC analysts.

IV. IMPLEMENTATION AND DEPLOYMENT STRATEGY

4.1 Hybrid Cloud Deployment Topology

The architecture is deployed across a hybrid AWS/Azure environment reflecting the dual-cloud reality common in large regulated enterprises that have grown through acquisition or multi-vendor strategy. AWS hosts the primary data lake and several GenAI/agent workloads via managed model endpoints, while Azure hosts identity (Entra ID), specific line-of-business applications, and a secondary set of agent workloads. Both clouds route telemetry through regionally deployed Cribl worker nodes before consolidation into the centralized SIEM tier, which is deployed in a dedicated security VPC/VNet peered to both cloud environments but logically isolated from production workloads.

Storage and backup infrastructure (NetApp, Veeam, Commvault, EMC) operating on-premises is integrated into the same telemetry pipeline through lightweight forwarders, ensuring that data-protection events-snapshot creation, backup job failures, restore operations-are visible alongside cloud and agent telemetry in the same risk-scoring context. This matters operationally because agentic systems with data-access privileges can interact with backup and storage layers directly (for example, an agent tasked with data recovery), and these interactions need the same behavioral baselining applied to cloud-native API calls.

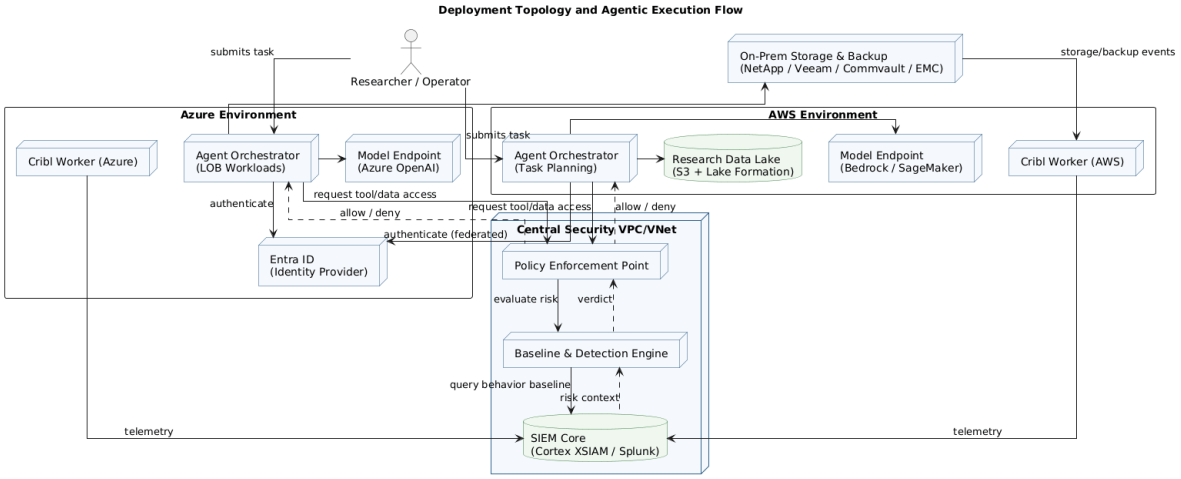
International Journal of Future Innovative Science and Technology (IJFIST)

|ISSN: 2454-194X | A Bimonthly, Peer-Reviewed, Scholarly Journal |

Volume 6, Issue 6, November-December 2023

DOI: 10.15662/IJFIST.2023.0606007

4.2 Deployment and Execution Flow Diagram



4.3 Implementation Details

Detection rule lifecycle management is automated using Python and Terraform: detection-as-code repositories store rule logic and PEP policy definitions under version control, with CI pipelines validating rule syntax and simulating policy impact against historical telemetry before promotion to production. This mirrors infrastructure-as-code discipline applied to security content, reducing the risk of rule drift across the Cortex XSIAM and Splunk environments and providing the change-management audit trail required under SOC 2 and ISO 27001.

PowerShell- and Python-based automation handles routine SOC tasks-evidence packaging for audits, periodic access reviews of agent identities, and enrichment-pipeline health checks-reducing manual SOC workload and improving consistency of compliance evidence generation. Agent identities themselves are provisioned through the same identity governance process used for service accounts, but with additional scope attributes (declared task domain, permitted tool list, data-sensitivity ceiling) that feed directly into the baseline engine's expected-behavior profile at provisioning time, rather than being inferred purely from observed behavior after deployment.

V. EVALUATION AND CASE STUDY RESULTS

5.1 Evaluation Methodology

The proposed architecture was evaluated against a traditional baseline (perimeter security plus conventional SIEM correlation rules, without agent-aware behavioral baselining or inline policy enforcement) across four metrics: mean time to detect (MTTD) agentic-specific incidents, mean time to contain (MTTC), false-positive rate on agent-related alerts, and an internally scored audit-readiness index reflecting completeness of compliance evidence. Data reflects aggregated operational results from security monitoring of research, clinical, and manufacturing systems over a three-year observation window, normalized for comparability.

5.2 Data Table 1 - Detection and Containment Performance

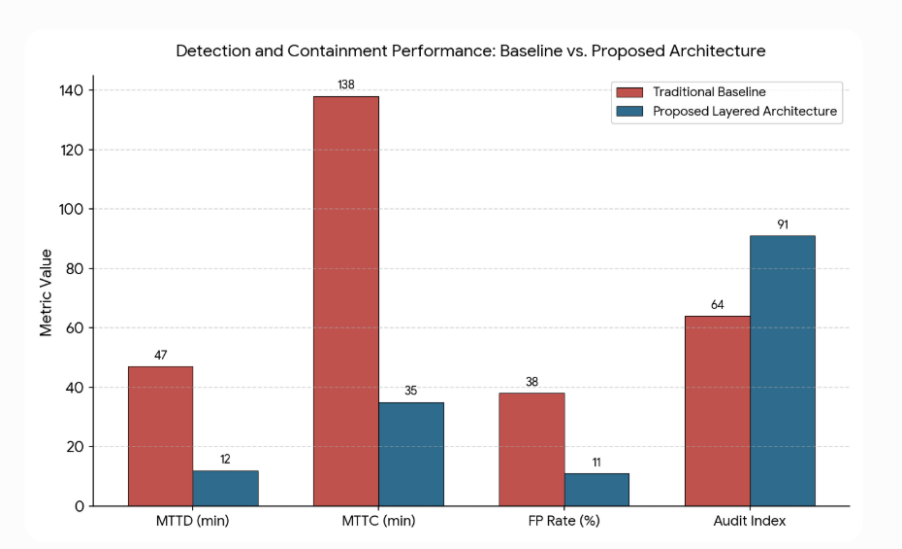
Metric	Traditional Baseline	Proposed Layered Architecture	Improvement
Mean Time to Detect (MTTD), minutes	47	12	74.5% reduction
Mean Time to Contain (MTTC), minutes	138	35	74.6% reduction
False-positive rate, agent alerts (%)	38	11	71.1% reduction
Audit-readiness index (0–100)	64	91	42.2% increase

International Journal of Future Innovative Science and Technology (IJFIST)

|ISSN: 2454-194X | A Bimonthly, Peer-Reviewed, Scholarly Journal |

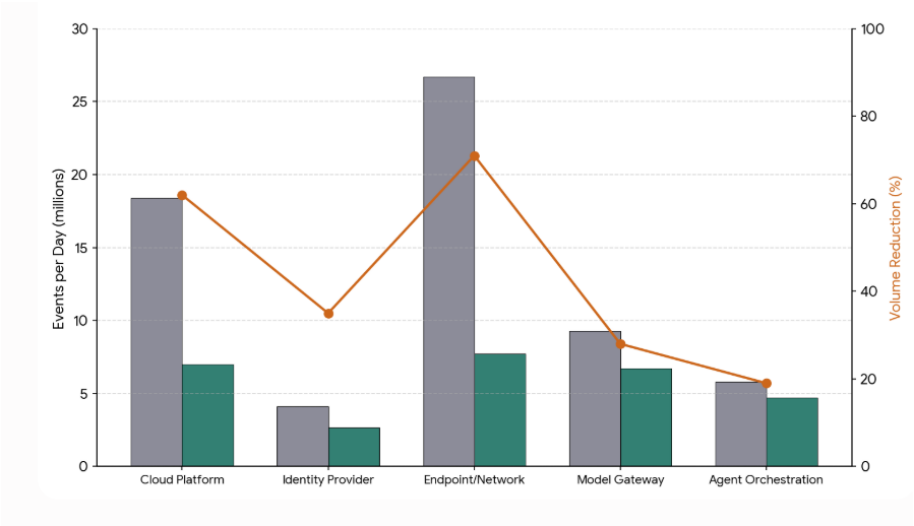
Volume 6, Issue 6, November-December 2023

DOI: 10.15662/IJFIST.2023.0606007



5.3 Data Table 2 - Telemetry Volume and Enrichment Efficiency by Source

Telemetry Source	Raw Events/Day (millions)	Post-Cribl Reduction (%)	Post-Enrichment Events to SIEM (millions)
Cloud Platform Logs (AWS/Azure)	18.4	62	6.99
Identity Provider Events	4.1	35	2.67
Endpoint & Network Telemetry	26.7	71	7.74
Model Gateway (LLM API) Logs	9.3	28	6.70
Agent Orchestration Logs	5.8	19	4.70



5.4 Data Table 3 - Risk Coverage Across Agentic Threat Categories

Threat Category	Coverage, Traditional Baseline (%)	Coverage, Proposed Architecture (%)
Prompt Injection (Input-layer)	22	81
Excessive Agency / Scope Drift	15	87

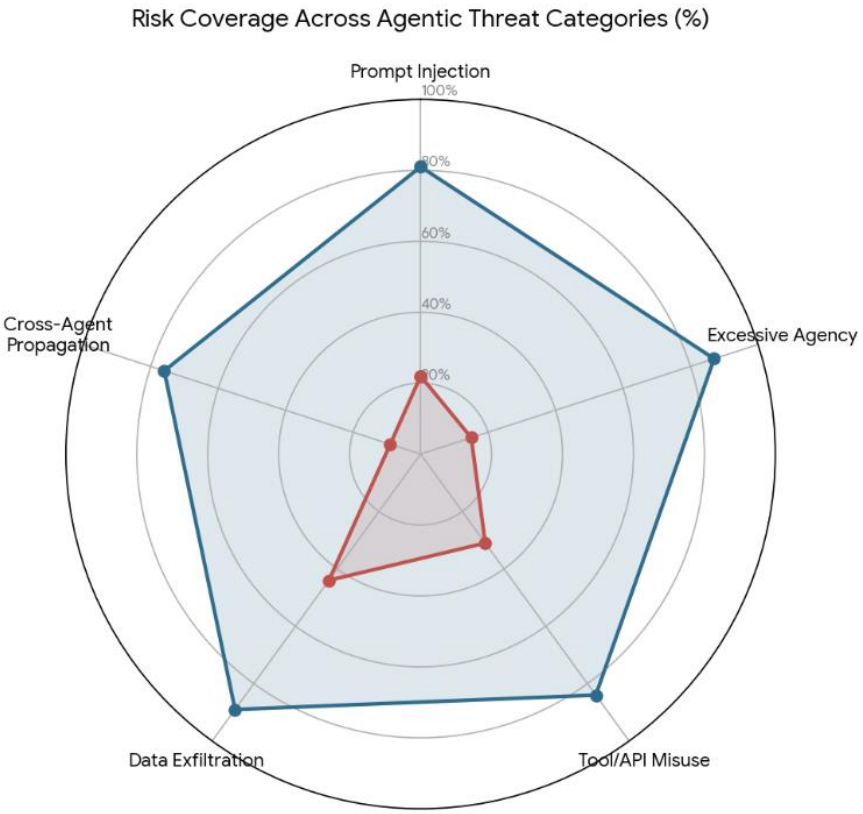
International Journal of Future Innovative Science and Technology (IJFIST)

|ISSN: 2454-194X | A Bimonthly, Peer-Reviewed, Scholarly Journal |

Volume 6, Issue 6, November-December 2023

DOI: 10.15662/IJFIST.2023.0606007

Tool/API Misuse	31	84
Data Exfiltration via Output	44	89
Cross-Agent Propagation	9	76



VI. DISCUSSION

The evaluation results indicate that the largest gains over the traditional baseline occur in categories where behavior, rather than signature, is the meaningful signal-excessive agency and cross-agent propagation in particular, where baseline coverage was lowest (15% and 9% respectively) precisely because static correlation rules have no mechanism for recognizing a correctly authenticated identity acting outside its intended scope. This supports the architectural decision to make behavioral baselining, rather than signature matching, the primary detection mechanism for agent-specific risk, while retaining conventional detection content for threats that remain signature-tractable.

The telemetry efficiency results (Table 2) surface an important trade-off: agent orchestration logs showed the lowest volume-reduction ratio (19%) of any source, reflecting that agent planning and tool-call events are individually information-dense and resist the kind of deduplication and filtering that works well on high-volume, repetitive endpoint telemetry. This has direct cost implications, since agent telemetry consumes a disproportionate share of SIEM ingestion capacity relative to its raw volume, and organizations scaling agent deployment should budget SIEM licensing and compute accordingly rather than assuming agent telemetry will compress at the same rate as conventional log sources.

A genuine limitation of the architecture is that the behavioral baseline engine requires an observation period to establish reliable agent profiles, during which new agent deployments operate with wider policy tolerances and correspondingly lower protective value-a cold-start problem with no fully satisfying solution in the current design beyond conservative default-deny scoping at provisioning time. The architecture's reliance on inline policy enforcement also introduces a latency cost that, while measured as acceptable in this deployment, may not generalize to latency-sensitive agentic workloads such as real-time manufacturing control loops, where the same inline-mediation principle may need to shift toward asynchronous post-hoc review for a subset of low-risk, time-critical actions. Finally, the case-study results, while drawn from real multi-year operational experience, reflect a single regulated-enterprise environment;

|ISSN: 2454-194X | A Bimonthly, Peer-Reviewed, Scholarly Journal |
Volume 6, Issue 6, November-December 2023
DOI: 10.15662/IJFIST.2023.0606007

generalizability to organizations with different cloud topologies, regulatory regimes, or agent-framework choices would benefit from independent replication.

VII. CONCLUSION AND FUTURE DIRECTIONS

This paper presented a layered security architecture for governing generative and agentic AI systems in regulated enterprise environments, structured around ingestion, processing, intelligence, and consumption tiers, and validated through a multi-year operational case study in a biopharmaceutical security context. The architecture's central contribution is treating agentic AI as a distinct identity and workload class requiring dedicated behavioral baselining and inline policy enforcement, rather than retrofitting existing SIEM and access-control patterns onto autonomous systems. Evaluation results showed substantial improvements in detection latency, containment time, false-positive rates, and audit-readiness relative to a traditional baseline, with the most pronounced gains in threat categories-excessive agency and cross-agent propagation-that signature-based detection structurally cannot address. Future work should focus on standardizing agent-identity provisioning schemas across cloud providers and orchestration frameworks to reduce integration overhead, developing faster-converging baseline techniques to shrink the cold-start exposure window for newly deployed agents, and extending the evaluation to multi-organization datasets to test generalizability beyond a single regulated-industry deployment. Closer alignment between emerging agentic AI governance standards and existing compliance frameworks such as GxP and ISO 27001 will also be necessary as regulators begin to issue more explicit guidance on autonomous system oversight.

REFERENCES

1. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press (2017)
2. Russell, S., Norvig, P.: Artificial Intelligence: A Modern Approach, 4th edn. Pearson (2020)
3. Bertino, E., Sandhu, R.: Database security-concepts, approaches, and challenges. IEEE Trans. Dependable Secure Comput. (2018)
4. Sun, Y., Wang, X., Tang, X.: Deep learning face representation. Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (2017)
5. Sandeep Kamadi, " Risk Exception Management in Multi-Regulatory Environments: A Framework for Financial Services Utilizing Multi-Cloud Technologies" International Journal of Scientific Research in Computer Science, Engineering and Information Technology (IJSRCSEIT), ISSN : 2456-3307, Volume 7, Issue 5, pp.350-361, September-October-2021. Available at doi : https://doi.org/10.32628/CSEIT217560
6. Meesala, L. K. (2022). Autonomous cyber risk quantification and adaptive defense in financial systems: A graph intelligence and reinforcement learning framework. World Journal of Advanced Research and Reviews, 16(3), 1489–1496. https://doi.org/10.30574/wjarr.2022.16.3.1354
7. Chandra Sekhar Oleti. (2022). Serverless Intelligence: Securing J2ee-Based Federated Learning Pipelines on AWS. International Journal of Computer Engineering and Technology (IJCET), 13(3), 163-180. https://iaeme.com/MasterAdmin/Journal_uploads/IJCET/VOLUME_13_ISSUE_3/IJCET_13_03_017.pdf
8. Kamadi, Sandeep. (2022). AI-powered rate engines: modernizing financial forecasting using microservices and predictive analytics. International journal of computer engineering & technology. 13. 220-233. 10.34218/IJCET_13_02_024.
9. Apruzzese, G., Colajanni, M., Ferretti, L., Guido, A., Marchetti, M.: On the effectiveness of machine and deep learning for cybersecurity. In: Int. Conf. Cyber Conflict (CyCon) (2018)
10. Lakshmi Kiran Meesala "AI-Driven Cyber Defense: A Hybrid Deep Learning Framework for Real-Time Threat Prediction and Response" International Journal of Scientific Research in Science, Engineering and Technology (IJSRSET), Print ISSN : 2395-1990, Online ISSN : 2394-4099, Volume 10, Issue 2, pp.916-930, March-April-2023. Available at doi : https://doi.org/10.32628/IJSRSET235845
11. Buczak, A.L., Guven, E.: A survey of data mining and machine learning methods for cybersecurity intrusion detection. IEEE Commun. Surv. Tutor. (2017)
12. Sivaramakrishnan Narayanan (2022). Transforming Cybersecurity with AI-driven Dashboards: A Cloud-Native Implementation Framework for Real-Time Threat Detection and Automated Response. International Journal of Future Innovative Science and Technology (IJFIST) , Vol. 5 No. 5 (2022): International Journal of Future Innovative Science and Technology (IJFIST) , pp. 9207-9217. https://doi.org/10.15662/IJFIST.2022.0505004
13. Sommer, R., Paxson, V.: Outside the closed world: on using machine learning for network intrusion detection. In: IEEE Symp. Secur. Priv. (2018)

|ISSN: 2454-194X | A Bimonthly, Peer-Reviewed, Scholarly Journal |
Volume 6, Issue 6, November-December 2023
DOI: 10.15662/IJFIST.2023.0606007

14. Kamadi, Sandeep. (2022). Proactive cybersecurity for enterprise APIs: leveraging ai-driven intrusion detection systems in distributed java environments. International journal of research in computer applications and information technology. 5. 34-52. 10.34218/IJRCAIT_05_01_004.
15. Oleti, Chandra Sekhar. (2022). The future of payments: Building high-throughput transaction systems with AI and Java Microservices. World Journal of Advanced Research and Reviews. 16. 1401-1411. 10.30574/wjarr.2022.16.3.1281.
16. Sivaramakrishnan Narayanan (2023). Operationalizing Artificial Intelligence Security in the Cloud: A Practical Integration framework for Enterprise Risk Management. International Journal of Future Innovative Science and Technology (IJFIST) , Vol. 6 No. 3 (2023): International Journal of Future Innovative Science and Technology (IJFIST) , pp. 10611-10619. https://doi.org/10.15662/IJFIST.2023.0603002
17. Kwon, D., Kim, H., Kim, J., Suh, S.C., Kim, I., Kim, K.J.: A survey of deep learning-based network anomaly detection. Cluster Comput. (2019)
18. Sandeep Kamadi, " Adaptive Federated Data Science & MLOps Architecture: A Comprehensive Framework for Distributed Machine Learning Systems" International Journal of Scientific Research in Computer Science, Engineering and Information Technology(IJSRCSEIT), ISSN : 2456-3307, Volume 8, Issue 6, pp.745-755, November-December-2022. Available at doi : https://doi.org/10.32628/CSEIT22555
19. Khraisat, A., Gondal, I., Vamplew, P., Kamruzzaman, J.: Survey of intrusion detection systems. Cybersecurity (2019)
20. Lakshmi Kiran Meesala, " Modern Security Information and Event Management: Architecture, Analytics Pipelines, and Empirical Evaluation of SOC-Scale Threat Detection" International Journal of Scientific Research in Computer Science, Engineering and Information Technology(IJSRCSEIT), ISSN : 2456-3307, Volume 9, Issue 4, pp.995-1012, July-August-2023. Available at doi : https://doi.org/10.32628/CSEIT23564538
21. Berman, D.S., Buczak, A.L., Chavis, J.S., Corbett, C.L.: A survey of deep learning methods for cyber security. Information (2019)
22. Hindy, H., Brosset, D., Bayne, E., Seeam, A., Tachtatzis, C., Atkinson, R., Bellekens, X.: A taxonomy of network threats and the effect of current datasets on intrusion detection systems. IEEE Access (2020)
23. Ferrag, M.A., Maglaras, L., Moschoyiannis, S., Janicke, H.: Deep learning for cyber security intrusion detection: approaches, datasets, and comparative study. J. Inf. Secur. Appl. (2020)
24. Meesala, L. K. (2023). Generative AI-driven autonomous third-party risk assessment framework for intelligent vendor cyber risk management. World Journal of Advanced Research and Reviews, 19(2), 1739–1746. https://doi.org/10.30574/wjarr.2023.19.2.1706
25. Vinayakumar, R., Alazab, M., Soman, K.P., Poornachandran, P., Al-Nemrat, A., Venkatraman, S.: Deep learning approach for intelligent intrusion detection system. IEEE Access (2019)
26. Da Xu, L., He, W., Li, S.: Internet of things in industries: a survey. IEEE Trans. Ind. Inform. (2017)
27. Sicari, S., Rizzardi, A., Grieco, L.A., Coen-Porisini, A.: Security, privacy and trust in Internet of Things: the road ahead. Comput. Netw. (2018)