Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

Bridging the Gap Between AI Innovation and Cybersecurity: Governance, Challenges, and Organizational Misalignment

Richard B. Antwi
University of the Cumberlands
Williamsburg, Kentucky

Abstract - The rapid deployment of artificial intelligence (AI) systems across organizations has shown a critical disconnect between the pace of technological innovation and the maturity of cybersecurity governance. While AI systems are increasingly embedded in high-stakes organizational processes, the implementation of cybersecurity controls remains hindered by persistent structural, technical, and organizational challenges. This qualitative study draws on semi-structured interviews with 12 AI and cybersecurity professionals across the financial services, healthcare, cloud computing, and technology sectors to examine the challenges of implementing AI cybersecurity controls and the role of AI governance in securing AI systems and sensitive data. Thematic analysis identified five major challenge domains — lack of standardized frameworks, AI development and cybersecurity team misalignment, tool and capability gaps, pace mismatch between innovation and governance, and regulatory ambiguity alongside four governance themes: formal AI governance structures and their limitations, compliance integration into the AI lifecycle, governance as a multiplier of security effectiveness, and the role of leadership accountability. The findings reveal that organizational and governance factors are at least as consequential as technical factors in determining AI security outcomes, and that bridging the gap between AI innovation and cybersecurity requires coordinated investment in governance, cross-functional collaboration, and implementation-level guidance.

Keywords — AI governance; cybersecurity challenges; organizational misalignment; AI security; qualitative research; risk management; regulatory compliance; DevSecOps; machine learning security; AI policy

I. INTRODUCTION

Artificial intelligence (AI) has become a transformative force across industries, enabling automation, predictive analytics, and intelligent decision-making at unprecedented scale. Financial institutions deploy machine learning (ML) models for real-time fraud detection; healthcare organizations use AI for clinical decision support; cloud providers offer AI-powered services that process billions of transactions daily [1]. Yet the security infrastructure supporting these deployments has not kept pace with the rate of adoption. The result is a growing gap between AI innovation and cybersecurity governance, one that creates systemic risk with potentially severe operational, regulatory, and reputational consequences [2].

This gap is not primarily a technical problem. Technical defenses for AI systems, such as adversarial training, data validation, differential privacy, and model monitoring, are well established in the research literature and are increasingly available in commercial tooling [3, 4]. The more fundamental challenge is organizational and structural: AI development teams and cybersecurity teams often operate with different priorities, different languages, and different incentive structures; governance frameworks lag behind deployment; and regulatory guidance for AI systems remains fragmented and ambiguous [5, 6].

Despite the urgency of these challenges, empirical research examining how practitioners experience and navigate them in real organizational settings remains limited. Most existing literature addresses AI security from a technical standpoint, proposing specific defenses for specific attack types [7, 8]. The organizational, governance, and strategic dimensions of AI security, which are the domain in which practitioners operate day-to-day, are substantially underexplored.

This paper addresses that gap by reporting findings from a qualitative study guided by two research questions: RQ1 — What challenges do AI professionals face in implementing cybersecurity controls in AI systems? and RQ2 — What role does AI governance play in securing AI systems and sensitive data? Drawing on semi-structured interviews with 12 practitioners, the study provides a practitioner-grounded analysis of the structural and organizational barriers to effective AI security, and of the role that governance frameworks, policies, and leadership play in either bridging or widening those gaps.

The findings have implications for practitioners designing AI security programs, for policymakers developing AI governance frameworks, and for researchers seeking to understand the sociotechnical dimensions of AI risk management.

II. LITERATURE REVIEW

A. The Organizational Dimensions of AI Security

The information security literature has long recognized that security is fundamentally a sociotechnical phenomenon: technical controls operate within organizational contexts that shape their implementation, enforcement, and effectiveness [9]. Studies of information security management have documented that security culture, governance structures, and cross-functional collaboration are as determinative of security outcomes as the quality of technical controls [10, 11]. Effective

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

security requires alignment between security functions and business operations, an alignment that is difficult to achieve and maintain amid rapid technological change [12].

In the AI security domain, these organizational dynamics are amplified by the novelty of the technology, the complexity of the systems, and the pace of deployment. Research on AI ethics and governance has documented persistent tensions between the imperatives of innovation and the requirements of responsible deployment [13]. Studies of AI in regulated industries have highlighted the particular challenges of operating under regulatory frameworks that were not designed with AI in mind [14]. Yet few studies have examined how practitioners experience these tensions and navigate them in their day-to-day work.

B. Challenges in AI Security Implementation

Several categories of implementation challenge have been identified in the AI security literature. The absence of standardized security frameworks specifically designed for AI systems has been widely noted as a structural gap that forces organizations to adapt conventional cybersecurity frameworks such as ISO/IEC 27001 [15] and the NIST Cybersecurity Framework [16] to AI-specific contexts for which they were not designed. While the NIST AI Risk Management Framework (AI RMF) [17] and ENISA's AI threat landscape [18] represent important steps toward AI-specific guidance, their adoption in practice and their translation into implementable controls remain limited.

The misalignment between AI development and cybersecurity teams has been identified as a particularly consequential structural challenge [19]. DevSecOps research has documented the difficulties of integrating security into fast-moving development pipelines [20], but AI development introduces additional complexity: ML systems have distinctive security properties including sensitivity to training data, susceptibility to adversarial manipulation, and emergent behavioral characteristics that require specialized expertise most cybersecurity teams do not possess [21].

Tool and capability gaps represent a further implementation challenge. Most commercial security tools were designed for traditional software systems and do not address AI-specific vulnerabilities such as model behavioral anomalies, adversarial inputs, or training data integrity [22]. This tooling gap forces organizations to build custom solutions or operate with inadequate visibility into AI-specific risk dimensions [23].

C. AI Governance: Frameworks, Policies, and Practice

AI governance has emerged as a major focus of policy, research, and practice as organizations and regulators grapple with the ethical, safety, and security implications of AI deployment. Governance frameworks, including the NIST AI RMF [17], the EU AI Act [24], and ISO/IEC 42001 [25], provide high-level principles and structured approaches to AI risk management. Model risk management (MRM) frameworks in financial services [26] represent a more established governance tradition that predates contemporary AI ethics discussions but offers relevant precedent for AI model oversight.

Despite the proliferation of governance frameworks, empirical evidence on their practical adoption and effectiveness is limited. Studies have documented significant gaps between governance policy development and implementation [27], and have noted that governance frameworks often lack the technical specificity needed to translate principles into operational controls [28]. Leadership commitment to AI governance has been identified as a critical enabler of effective implementation, but the mechanisms through which leadership accountability shapes security outcomes remain underexplored [29].

III. RESEARCH METHODOLOGY

A. Research Design

A qualitative research design employing semi-structured interviews was used to explore the organizational and governance dimensions of AI security. This approach was selected as appropriate for the study's purpose of examining complex, context-dependent practitioner experiences that cannot be adequately captured through quantitative methods [30]. Semi-structured interviews provided the flexibility to pursue emergent themes while maintaining sufficient consistency across participants for systematic comparative analysis [31].

B. Sampling and Participants

Twelve AI and cybersecurity professionals were recruited through purposive sampling to ensure representation of diverse functional roles, organizational contexts, and industry sectors relevant to AI security. Eligibility required substantive professional involvement in the development, deployment, or governance of AI systems. Snowball sampling supplemented initial recruitment. Table I summarizes participant profiles.

TABLE I. Participant Profiles

ID	Role	Exp.	Sector
P1	Senior AI Security & Risk Analyst	7+ yr	Financial Services
P2	Lead ML Engineer (AI Risk)	10+ yr	FinTech/Digital Banking
P3	Senior AI Security Engineer / MLOps	9+ yr	Cloud AI / SaaS
P4	Director, AI Governance & Compliance	12+ yr	Banking
P5	Senior Data Scientist	8+ yr	Healthcare Analytics
P6	AI Product Lead / Co-Founder	6+ yr	AI SaaS Startup
P7	Senior IT Auditor / AI Risk Auditor	11+ yr	Banking
P8	Cloud Security Architect (AI/ML)	10+ yr	Enterprise Cloud
P9	AI Privacy & Responsible AI Lead	9+ yr	Global Technology
P10	Senior AI Security Consultant	13+ yr	Consulting (Multi)
P11	Senior DevSecOps Engineer (AI/ML)	9+ yr	Enterprise Tech
P12	CISO	18+ yr	Financial Services

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

C. Data Collection

Semi-structured interviews were conducted using a protocol organized into six thematic sections. Questions relevant to RQ1 explored the specific challenges participants face in implementing AI security controls, including structural barriers, tool limitations, team collaboration difficulties, and regulatory compliance challenges. Questions relevant to RQ2 examined the governance structures, policies, and leadership mechanisms used to manage AI security risk. Interviews were conducted virtually, audio-recorded with participant consent, and professionally transcribed. The duration ranged from 45 to 75 minutes.

D. Data Analysis

Thematic analysis was conducted following the six-phase process described by Braun and Clarke [32]: data familiarization, initial coding, theme construction, theme review, theme definition, and report production. Coding proceeded inductively from the data. Codes were iteratively refined and clustered into higher-order themes reflecting dominant patterns across participant accounts. Analytical rigor was supported through member checking, reflexive journaling, and peer debriefing. Thematic saturation was observed after approximately ten interviews.

IV. FINDINGS: IMPLEMENTATION CHALLENGES (RQ1)

Thematic analysis of participant accounts identified five major challenge domains in implementing AI cybersecurity controls. These challenges are interrelated and mutually reinforcing, creating complex barriers that cannot be resolved through isolated technical or organizational interventions.

A. Challenge 1: Absence of Standardized AI Security Frameworks

The most frequently cited structural challenge was the absence of standardized, implementation-ready security frameworks specifically designed for AI systems. Participants across all roles and sectors described the resulting organizational uncertainty and the burden of developing custom approaches in the absence of clear guidance.

P1 identified this as a primary barrier: “One major challenge is the lack of standardized security practices for AI systems.” P2 elaborated on the consequences: “One of the biggest challenges is the lack of standardized practices for AI security. Unlike traditional cybersecurity, where frameworks are well-established, AI security is still evolving.” P10 observed the same pattern across consulting clients: “One of the biggest challenges is the lack of standardization. Organizations are experimenting with different approaches, and there is no universally accepted framework for AI security.”

The NIST AI RMF was frequently acknowledged as a valuable reference but consistently described as insufficiently specific for operational implementation. P3 noted: “Frameworks provide good guidance, but they can be difficult to implement in practice, especially in fast-paced environments. There is a clear need for tools and frameworks that bridge the gap between high-level guidance and hands-on implementation.” P4 elaborated: “Many frameworks provide high-level guidance but lack practical implementation details. This can make it difficult for organizations to translate governance requirements into actionable controls.” P11 argued that framework effectiveness depends entirely on technical operationalization: “Frameworks are useful, but they need to be implemented through automation to be effective. A policy that is not enforced through automation is difficult to sustain in a high-speed environment.”

B. Challenge 2: AI Development and Cybersecurity Team Misalignment

All 12 participants described a persistent structural gap between AI development teams and cybersecurity teams as a significant and consequential implementation challenge. This misalignment was attributed to fundamental differences in professional priorities, technical expertise, organizational incentives, and cultural norms.

P1 identified the core dynamic: “AI teams focus on performance, while cybersecurity focuses on risk. Misalignment often occurs during development and deployment stages. This can lead to security being overlooked initially.” P2 described the organizational consequences: “AI teams are primarily focused on model performance and innovation, while cybersecurity teams are focused on risk and control. These priorities don't always align. This can lead to situations where security considerations are introduced late in the development process.”

P7, operating from an independent audit perspective, characterized this misalignment as a systematic governance failure: “AI development teams prioritize performance and innovation, while cybersecurity teams focus on risk mitigation. This misalignment can result in security controls being applied after the fact, rather than being integrated into the development process from the beginning.” P10 identified the timing consequence as particularly costly: “I often see security being introduced too late in the development process, which makes it harder and more expensive to implement effectively.”

P3 described a structural response being explored in more advanced organizations: “We are trying to address this by embedding security engineers within AI teams, but this is still not standard practice.” P4 described a governance-level intervention: “We are addressing this by establishing cross-functional governance committees and encouraging collaboration, but it remains an ongoing challenge.” P9 highlighted the multi-team dimension: “There are also gaps between AI teams and privacy teams. Each group has a different focus. AI teams prioritize performance, cybersecurity teams focus on threats, and privacy teams focus on data protection and compliance. These differences can lead to fragmented approaches to risk management if there is no strong governance structure to align them.”

C. Challenge 3: Tool and Capability Gaps

Participants consistently described existing security tools as designed for traditional IT systems and inadequately adapted to the distinctive characteristics of AI systems. This tool gap forces organizations to build custom solutions, rely on

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

incomplete capabilities, or accept residual risk they cannot effectively monitor or mitigate.

P3 described the technical specificity of the gap: "Most existing tools are not designed with AI in mind. For example, traditional vulnerability scanners do not assess model vulnerabilities. Frameworks may recommend 'secure the model,' but they don't specify how to protect against model extraction or adversarial attacks." P11 identified a parallel limitation in DevSecOps tooling: "Many tools are designed for traditional software development and do not fully address the unique aspects of AI systems. For example, tools can scan code and containers, but they may not assess model behavior or data-related risks."

P1 characterized the tool situation broadly: "Current tools are not fully designed for AI-specific risks. Many tools are adapted from traditional cybersecurity. Features like model monitoring and adversarial testing are limited." P10 observed this pattern across client organizations: "Most tools are not designed specifically for AI systems, which creates limitations in addressing AI-specific risks. There is a clear need for tools and frameworks that bridge the gap between high-level guidance and hands-on implementation."

A specific capability gap identified by multiple participants was the absence of production-ready tools for monitoring model behavior. P7 described this as a consistent audit finding: "A gap I frequently observe is that organizations monitor system uptime and performance very well, but they do not always have sufficient controls in place to monitor the integrity and behavior of the AI models themselves." P8 noted the correlation challenge: "One of the key challenges is correlating signals across different layers. What looks like a normal system event might actually be part of a larger attack when viewed in context."

D. Challenge 4: Pace Mismatch Between Innovation and Security

A fourth major challenge was the structural tension between the speed of AI development and deployment and the slower pace at which security controls and governance frameworks can be developed, implemented, and maintained. Participants described this as a fundamental organizational dynamic that creates persistent security gaps.

P2 described the development velocity pressure: "Another challenge is the speed at which AI systems are developed and deployed. There is often pressure to deliver quickly, which can lead to security being treated as a secondary concern." P4 characterized this from a governance perspective: "One of the biggest challenges is keeping up with the pace of AI innovation. New models and use cases are being developed rapidly, often faster than governance frameworks can adapt."

P6, representing a startup context, described the manifestation of this challenge under resource constraints: "The biggest challenge is balancing speed and security. In a startup, there is constant pressure to move fast, release features, and stay competitive. Security can sometimes feel like a constraint, especially when resources are limited." P12 identified the executive-level governance implication: "AI is often seen as a driver of innovation and competitive advantage, which can create pressure to move quickly. At the same time, we must ensure that risks are properly managed, which requires careful governance and oversight. Balancing innovation and risk is one of the most complex challenges at the executive level."

P11 proposed automation as the primary mechanism for resolving the pace tension: "Security needs to move at the same speed as development, otherwise it becomes a bottleneck rather than an enabler. Balancing speed, automation, and security is one of the hardest problems in DevSecOps for AI systems."

E. Challenge 5: Regulatory Ambiguity and Compliance Complexity

A fifth challenge concerned the complexity and ambiguity of the regulatory environment for AI systems. Participants described the dual pressure of compliance with existing regulations that were not designed for AI and the uncertainty created by rapidly evolving AI-specific regulatory frameworks.

P1 identified the core interpretive challenge: "Challenges include interpreting regulations for AI-specific use cases." P4 elaborated: "Many regulations were not designed specifically for AI, so we often have to interpret how they apply to our systems. Challenges include interpreting regulations for AI-specific use cases. Regulations sometimes slow down innovation but are necessary." P7 noted the regulatory evolution dynamic: "As AI systems evolve, regulatory expectations are also changing, which creates uncertainty. One of the challenges is that many regulations were not designed with AI in mind, which requires organizations to interpret how these requirements apply to AI systems."

P9 identified a governance-level consequence of regulatory ambiguity: "One of the challenges is the lack of clear regulatory guidance for AI systems. While data protection laws exist, they do not always address the nuances of AI. Organizations are often navigating uncharted territory when it comes to balancing innovation with responsibility." P12 described the organizational imperative: "Organizations must be proactive in interpreting and applying regulations to AI systems, rather than waiting for explicit guidance."

Table II summarizes the five challenge domains, their primary manifestations, and representative organizational responses identified across participant accounts.

TABLE II. AI Security Implementation Challenges: Domains and Manifestations

Challenge Domain	Primary Manifestations	Representative Response
Absence of Standardized Frameworks	No consensus AI security standards; high-level frameworks lack technical specificity; organizations develop custom approaches	Adapt NIST AI RMF; build internal standards; invest in specialized expertise
Team Misalignment	Divergent priorities between AI dev and security teams; security introduced late; siloed expertise	Cross-functional governance committees; embedded security engineers;

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

		DevSecOps adoption
Tool & Capability Gaps	Traditional tools inadequate for AI risks; no standard model monitoring tools; limited adversarial testing tooling	Custom solution development; partial tool adaptation; manual processes
Pace Mismatch	AI deployment outpaces governance; development velocity pressure; framework evolution lag	Security automation; pipeline integration; risk appetite governance
Regulatory Ambiguity	Regulations not designed for AI; interpretive burden on practitioners; evolving requirements	Conservative compliance interpretation; legal team engagement; proactive regulatory monitoring

V. FINDINGS: AI GOVERNANCE ROLE (RQ2)

Thematic analysis of participant accounts regarding AI governance identified four major themes describing how governance structures, policies, frameworks, and leadership shape AI security outcomes.

A. Theme 1: Formal Governance Structures and Their Limitations

All participants confirmed that their organizations had some form of AI governance structure in place. However, the maturity, comprehensiveness, and operational effectiveness of these structures varied substantially, and participants consistently identified gaps between governance policy and implementation.

P1 described the state of governance evolution: “Yes, but they are still evolving. Policies cover data usage, model validation, and risk assessment. Governance is enforced by risk and compliance teams. Policies are becoming more formalized over time.” P4, whose role centered on AI governance, described a more formal structure: “We have established formal AI governance policies that cover model development, validation, deployment, and monitoring. These policies define roles and responsibilities, required controls, and approval processes. However, governance is not static, it is continuously evolving as we learn more about AI risks.”

P7, from an audit perspective, identified inconsistency of implementation as the primary governance gap: “Most organizations I work with have some form of AI governance policy, particularly in regulated industries. However, the level of maturity varies, and in some cases, policies are still evolving. Having a policy is one thing, but ensuring consistent implementation and enforcement across the organization is where many challenges arise.” P10 observed a common pattern across client organizations: “Most organizations I work with are in the process of developing AI governance policies, but the level of maturity varies. A common issue is that governance policies exist, but they are not fully integrated into day-to-day operations.”

P6, representing a startup context, described an early-stage governance situation: “We have some internal guidelines, but I wouldn't describe them as fully developed governance policies yet. Governance is something we are actively working on, especially as we engage with larger clients who expect more formal controls.”

B. Theme 2: Compliance Integration into the AI Lifecycle

A second governance theme concerned how regulatory compliance requirements are integrated into the AI development and deployment lifecycle. Participants described compliance as a governance function that spans multiple stages of the AI lifecycle and requires active interpretation and operationalization.

P1 described the integration approach: “Compliance is integrated into the development lifecycle. We ensure compliance through audits and controls. Regular assessments are conducted.” P2 noted the financial services regulatory context: “Compliance is taken very seriously, especially in financial services. We align with data protection regulations and internal policies, and we conduct regular audits to ensure adherence. One challenge is that many regulations were not designed specifically for AI, so we often have to interpret how they apply to our systems.”

P4 described a comprehensive lifecycle integration approach: “Compliance is integrated into every stage of the AI lifecycle. We conduct risk assessments, audits, and reviews to ensure that systems meet regulatory requirements. We align with financial regulations and data protection laws, but one challenge is that many regulations do not explicitly address AI.” P9, focused on privacy compliance, described privacy impact assessments as a key governance mechanism: “Privacy impact assessments are a key part of our process, and we work closely with legal and compliance teams to interpret regulatory requirements. One of the challenges is that regulations are evolving, and organizations must be proactive in adapting to new expectations.”

P11 described the challenge of operationalizing compliance through technical controls: "Compliance requirements are integrated into pipelines through automated checks and controls. For example, we enforce validation steps and audit logging to ensure traceability. However, interpreting regulatory requirements in the context of AI systems can be challenging."

C. Theme 3: Governance as a Multiplier of Security Effectiveness

A third governance theme, which emerged strongly across participant accounts, was the role of governance maturity as a multiplier of technical security control effectiveness. Participants consistently described the same technical controls as producing substantially different security outcomes depending on the quality of the governance context in which they were implemented.

P4 articulated this principle explicitly: “From a governance standpoint, we require that all AI systems adhere to enterprise cybersecurity standards. What is particularly important in AI systems is the extension of these controls into the model lifecycle. Controls must be documented, enforced, and auditable. AI security is not just a technical issue; it is a

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

governance issue." P12 framed governance as a necessary enabling condition for effective security at the enterprise level: "Governance is essential because it provides the structure and accountability needed to manage AI risks effectively."

P7 identified governance maturity as the primary differentiator of security effectiveness across audited organizations: "From an audit perspective, frameworks are valuable because they provide structure, but their effectiveness depends on how well they are implemented. Organizations are often compliant with existing standards, but compliance does not necessarily equate to security, especially in the context of AI systems." P10 observed a maturity-effectiveness correlation across consulting clients: "Effectiveness is highly dependent on organizational maturity. In more mature environments, security measures are integrated into the AI lifecycle. In less mature environments, security is often reactive."

P9 connected governance quality to trust outcomes: "Effective governance requires not only policies, but also strong oversight and a culture of accountability. Organizations that fail to prioritize privacy and security in their AI systems risk not only regulatory penalties, but also loss of user trust, which can be far more damaging in the long term."

D. Theme 4: Leadership Accountability and Governance Culture

A fourth governance theme concerned the role of leadership commitment and organizational culture in enabling effective AI governance. Participants described executive-level support, clear accountability structures, and a culture that treats security as a shared organizational responsibility as critical enablers of governance effectiveness.

P4 emphasized the cultural dimension: "Effectiveness depends on enforcement and organizational culture. Policies alone are not enough; they must be supported by strong oversight and accountability." P12, as CISO, described the executive governance responsibility: "At the executive level, my responsibility is not just technical security, but ensuring that the organization understands and manages risk appropriately. This includes setting risk appetite, aligning security with business objectives, and providing assurance to the board and regulators. AI security must become a strategic priority, not just a technical concern."

P2 identified leadership investment as a prerequisite for closing the gap between awareness and implementation: "Leadership also needs to take AI risk more seriously and invest in the necessary resources." P10 described the cultural transformation required: "AI security is not a one-time effort, it is an ongoing process. Organizations that treat AI security as a continuous journey, rather than a one-time implementation, are far more likely to succeed in managing risks effectively."

P11 described the governance culture from a DevSecOps perspective: "A policy that is not enforced through automation is difficult to sustain in a high-speed environment. In modern AI environments, the pipeline is the system; if you secure the pipeline, you significantly reduce the overall risk. Security must be embedded into the lifecycle, not added afterward." P1 offered a forward-looking assessment: "AI security will become more structured and regulated. New threats like advanced adversarial attacks will emerge. Organizations will need stronger governance. Technologies like explainable AI and automated monitoring will play a key role."

Table III summarizes the governance themes and their key mechanisms as identified across participant accounts.

TABLE III. AI Governance Themes and Key Mechanisms

Governance Theme	Key Mechanisms	Critical Gaps Identified
Formal Governance Structures	AI governance policies; model risk management; pre-deployment review processes; cross-functional governance committees	Policy-implementation gap; inconsistent enforcement; maturity variation across organizations
Compliance Integration	Risk assessments; audit cycles; privacy impact assessments; automated compliance controls in pipelines	Regulatory ambiguity for AI; interpretation burden; evolving requirements outpacing policies
Governance as Security Multiplier	Documentation and auditability; accountable oversight; governance-embedded technical controls	High-level frameworks lack technical specificity; compliance without security outcomes
Leadership & Culture	Executive risk ownership; board-level AI risk reporting; security-as-shared-responsibility culture; continuous improvement	Reactive rather than proactive posture; insufficient leadership prioritization of AI risk

VI. DISCUSSION

A. The Primacy of Organizational Over Technical Challenges

The findings of this study suggest that the primary barriers to effective AI security are organizational and structural rather than technical. The five challenge domains identified are the absence of a framework, team misalignment, tool gaps, pace mismatch, and regulatory ambiguity. All have technical dimensions, but they are rooted in organizational dynamics that technical fixes alone cannot resolve. This finding extends the broader information security literature's recognition that security is a sociotechnical problem [9, 10] to the specific context of AI systems, where the innovation and complexity of the technology amplifies the importance of organizational and governance factors.

The persistence of the AI development-cybersecurity team gap across all 12 participant accounts, across diverse organizational contexts and roles, suggests that this misalignment is a systemic characteristic of current AI deployment practice rather than an idiosyncratic organizational failure. Structural interventions such as embedded security

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

engineers, cross-functional governance committees, shared accountability frameworks, and DevSecOps integration were described as promising responses but not yet standard practice. Accelerating the adoption of these structural responses represents an important priority for both organizational practice and professional standards development.

B. The Framework Gap: From Principles to Implementation

The consistent characterization of existing AI governance frameworks as providing valuable high-level principles but insufficient operational guidance points to a critical gap in the current AI security landscape. The NIST AI RMF [17], ENISA threat landscape [18], and ISO/IEC 42001 [25] were all referenced as useful starting points that nonetheless require substantial organizational interpretation and customization before they can be operationalized as technical controls.

This framework gap has direct consequences for security outcomes. When practitioners must develop their own implementation approaches in the absence of clear guidance, the result is inconsistency across organizations, excessive burden on individual teams, and systematic underinvestment in areas such as adversarial testing and model behavioral monitoring, where the implementation path is least clear. Closing the framework-to-implementation gap through more technically specific, tiered, and sector-adapted guidance is an important priority for standards bodies and the research community.

C. Governance as Infrastructure, Not Overhead

Perhaps the most policy-relevant finding of this study is the characterization of AI governance not as compliance overhead but as a critical enabling infrastructure for security effectiveness. Participants consistently described organizations with mature governance structures as achieving substantially better security outcomes than those relying on technical controls without governance support. This finding has important implications for how AI security investment should be framed at the executive and board level.

If governance maturity multiplies the effectiveness of technical controls, then investment in governance in formal policy development, cross-functional oversight structures, leadership accountability mechanisms, and culture change generates security returns that extend across all technical security investments. Conversely, organizations that invest heavily in technical controls without corresponding governance development may find that those controls underperform relative to their potential. This multiplier dynamic suggests that AI security investment strategies should explicitly include governance development alongside technical capability building.

D. The Pace-Security Tension and Automation as Resolution

The tension between AI development velocity and security governance pace described by participants is a structural feature of the current AI deployment environment that is unlikely to resolve without deliberate organizational and technical intervention. The principle articulated by P11 that security needs to move at the same speed as development points to automation as the primary mechanism for bridging this pace gap.

The DevSecOps literature supports this principle, documenting that security automation integrated into continuous integration and delivery (CI/CD) pipelines can maintain security enforcement without creating development bottlenecks [20]. In the AI context, this principle extends to ML pipelines: automated model validation, automated security scanning of training data, automated deployment approval gates, and automated model behavioral monitoring represent the technical expression of governance requirements in high-velocity environments. The consistent observation that manual security processes are fragile in fast-moving AI development contexts suggests that automation is not merely an efficiency consideration but a prerequisite for reliable security.

E. Regulatory Evolution and Proactive Compliance

The regulatory ambiguity described by participants reflects a broader challenge facing AI governance: the gap between the pace of AI deployment and the pace of regulatory adaptation. Major regulatory developments, including the EU AI Act [24] and sector-specific guidance from financial regulators, represent important steps toward AI-specific regulatory clarity. However, participants' descriptions of operating under existing regulations that were not designed for AI suggest that the interpretive burden on practitioners will remain substantial for the foreseeable future.

The participants' call for proactive rather than reactive engagement with regulatory requirements, that is, interpreting and applying existing regulations to AI contexts without waiting for explicit guidance, points to a governance posture that requires substantial legal, technical, and compliance expertise working in close coordination. Developing this cross-disciplinary governance capability represents an important organizational investment for any organization with significant AI deployments in regulated sectors.

VII. CONCLUSION

This study examined the challenges AI professionals face in implementing cybersecurity controls in AI systems (RQ1) and the role of AI governance in securing those systems (RQ2), drawing on semi-structured interviews with 12 practitioners across diverse organizational contexts. Five major challenge domains were identified: absence of standardized AI security frameworks, AI development-cybersecurity team misalignment, tool and capability gaps, pace mismatch between innovation and governance, and regulatory ambiguity. Four governance themes were identified: formal governance structures and their limitations, compliance integration into the AI lifecycle, governance as a multiplier of technical security effectiveness, and the role of leadership accountability and organizational culture.

Together, these findings demonstrate that effective AI security is fundamentally an organizational and governance challenge as much as a technical one. The most consequential barriers to effective AI security are team misalignment, framework gaps, pace mismatch, and governance immaturity, which cannot be resolved through technical investment alone. They require structural organizational interventions, including

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

cross-functional governance mechanisms, embedded security expertise within AI teams, leadership accountability for AI risk, and significant improvement in the implementation-level specificity of AI security frameworks.

For practitioners, the findings underscore the imperative of treating AI governance as critical security infrastructure rather than compliance overhead, investing in cross-functional collaboration structures that bridge the AI development-security team gap, and prioritizing security automation as the primary mechanism for maintaining security pace with AI deployment velocity. For policymakers and framework developers, the consistent identification of implementation guidance deficits points to a critical need for more technically specific, tiered, and sector-adapted AI security guidance that translates high-level principles into operational controls.

Limitations of this study include its qualitative design and its relatively small purposive sample, which is concentrated in the financial services and technology sectors. Future research should examine these dynamics through larger quantitative samples, evaluate the effectiveness of specific governance interventions through longitudinal study designs, and examine how AI security governance challenges and practices vary across national regulatory environments and organizational maturity levels.

ACKNOWLEDGMENT

The author sincerely thanks all 12 research participants for their time and practitioner insights, which form the empirical foundation of this work.

REFERENCES

[1] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed. Upper Saddle River, NJ: Pearson, 2020.
[2] M. Taddeo, T. McCutcheon, and L. Floridi, "Trusting artificial intelligence in cybersecurity is a double-edged sword," Nature Machine Intelligence, vol. 1, pp. 557–560, 2019.
[3] B. Biggio and F. Roli, "Wild patterns: Ten years after the rise of adversarial machine learning," Pattern Recognition, vol. 84, pp. 317–331, 2018.
[4] M. Goldblum et al., "Dataset security for machine learning: Data poisoning, backdoor attacks, and defenses," IEEE Trans. Pattern Anal. Mach. Intell., vol. 45, pp. 1563–1580, 2022.
[5] A. Jobin, M. Ienca, and E. Vayena, "The global landscape of AI ethics guidelines," Nature Machine Intelligence, vol. 1, pp. 389–399, 2019.
[6] P. Cihon, J. Schuett, and D. Hadfield-Menell, "AI standards in industry: Understanding the landscape and governance challenges," SSRN, 2023.
[7] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in Proc. ICLR, San Diego, CA, 2015.
[8] N. Carlini and D. Wagner, "Towards evaluating the robustness of neural networks," in Proc. IEEE Symp. Security Privacy, San Jose, CA, 2017, pp. 39–57.
[9] Z. A. Soomro, M. H. Shah, and J. Ahmed, "Information security management needs more holistic approach: A literature review," Int. J. Inf. Manage., vol. 36, pp. 215–225, 2016.
[10] H. A. Kruger and W. D. Kearney, "A prototype for assessing information security awareness," Computers & Security, vol. 25, pp. 289–296, 2006.
[11] R. Von Solms and J. Van Niekerk, "From information security to cyber security," Computers & Security, vol. 38, pp. 97–102, 2013.
[12] K. J. Knapp, T. E. Marshall, R. K. Rainer, and F. N. Ford, "Information security: Management's effect on culture and policy," Inf. Manage. Comput. Security, vol. 14, pp. 24–36, 2006.
[13] M. Coeckelbergh, AI Ethics. Cambridge, MA: MIT Press, 2020.
[14] E. J. Topol, "High-performance medicine: The convergence of human and artificial intelligence," Nature Medicine, vol. 25, pp. 44–56, 2019.
[15] International Organization for Standardization, ISO/IEC 27001:2022: Information Security Management Systems—Requirements. Geneva: ISO, 2022.
[16] National Institute of Standards and Technology, Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1. Gaithersburg, MD: NIST, 2018.
[17] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0). Gaithersburg, MD: NIST, 2023. [Online]. Available: https://airc.nist.gov/RMF
[18] European Union Agency for Cybersecurity (ENISA), ENISA Threat Landscape for Artificial Intelligence. Athens: ENISA, 2023.
[19] S. Ransbotham and D. Kiron, "Analytics as a source of business innovation," MIT Sloan Management Review, vol. 58, pp. 1–6, 2017.
[20] R. N. Rajapakse, M. Zahedi, M. A. Babar, and H. Shen, "Challenges and solutions when adopting DevSecOps: A systematic review," Inf. Softw. Technol., vol. 141, p. 106700, 2022.
[21] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, "Practical black-box attacks against machine learning," in Proc. ACM Asia Conf. Comput. Commun. Security, Abu Dhabi, 2017, pp. 506–519.
[22] D. S. Berman, A. L. Buczak, J. S. Chavis, and C. L. Corbett, "A survey of deep learning methods for cyber security," Information, vol. 10, p. 122, 2019.
[23] A. Chakraborty, M. Alam, V. Dey, A. Chattopadhyay, and D. Mukhopadhyay, "Adversarial attacks and defences: A survey," arXiv:1810.00069, 2018.
[24] European Parliament, Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act). Brussels: European Parliament, 2024.
[25] International Organization for Standardization, ISO/IEC 42001:2023: Artificial Intelligence—Management System. Geneva: ISO, 2023.
[26] Board of Governors of the Federal Reserve System, SR 11-7: Guidance on Model Risk Management. Washington, DC: Federal Reserve, 2011.
[27] I. D. Raji et al., "Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing," in Proc. ACM FAccT, Barcelona, 2020, pp. 33–44.
[28] P. Cihon, "Standards for AI governance: International standards to enable global coordination in AI research and development," Future of Humanity Institute, Technical Report, 2019.
[29] L. Floridi and M. Chiriatti, "GPT-3: Its nature, scope, limits, and consequences," Minds and Machines, vol. 30, pp. 681–694, 2020.
[30] J. W. Creswell and J. D. Creswell, Research Design: Qualitative, Quantitative, and Mixed Methods Approaches, 5th ed. Thousand Oaks, CA: SAGE Publications, 2018.
[31] S. B. Merriam and E. J. Tisdell, Qualitative Research: A Guide to Design and Implementation, 4th ed. San Francisco, CA: Jossey-Bass, 2015.
[32] V. Braun and V. Clarke, "Using thematic analysis in psychology," Qualitative Research in Psychology, vol. 3, pp. 77–101, 2006.