

DOI: https://doi.org/10.48009/2_iis_2026_129

An AI–Cybersecurity Taxonomy for Regulating Interconnected Risks

Angelica Marotta, *MIT Sloan School of Management*, amarotta@mit.edu

Stuart Madnick, *MIT Sloan School of Management*, smadnick@mit.edu

Abstract

This paper investigates the need for increased coordination of AI and cybersecurity regulation, challenging the longstanding treatment of these fields as separate governance domains. It presents the AI–Cybersecurity Taxonomy, which identifies five core governance clusters that capture how regulatory frameworks are adapting cybersecurity governance to AI-specific security challenges. The study introduces key concepts, such as provisional governance, agentic amplification, and embedded verifiability, to explain why coordination across AI and cybersecurity obligations is becoming increasingly important. Through sector-specific analysis, it underscores the importance of flexible, layered governance models. The paper concludes with practical recommendations to better align regulation with the operational realities of AI-enabled cybersecurity risk.

Keywords: artificial intelligence, cybersecurity, regulation, taxonomy, compliance

Introduction

Regulatory frameworks are conventionally structured around distinct policy domains, with technologies classified according to the risks they present and governed through specialized instruments tailored to address those risks. This categorical approach is evident in how two critical technology areas have been regulated. Artificial intelligence (AI) is generally framed as a matter of algorithmic accountability, with regulatory attention directed toward fairness, transparency, and decision-making processes. For instance, the European Union’s AI Act emphasizes risk classification, conformity assessments, and transparency obligations for high-risk AI systems (European Union, 2024). Cybersecurity regulation, by contrast, is usually centered on system protection, prioritizing access control, data integrity, and threat mitigation. The Network and Information Security Directive 2 (NIS2 Directive), for example, focuses on incident reporting, security measures, and supply chain risk management for essential and important entities (European Union, 2022). This regulatory separation rests on a foundational assumption: Cybersecurity emerged first as a distinct field, followed by AI, leading to their governance through parallel regulatory frameworks with minimal coordination or cross-domain consideration, each addressing unique categories of risk (Marotta & Madnick, 2021).

This assumption is increasingly challenged by the pervasive integration of AI within critical operational infrastructure systems. According to Cisco’s 2026 State of Industrial AI Report, 61% of industrial organizations have moved AI into active, scaled deployment in live operational environments (Cisco, 2026). Therefore, contemporary AI systems are no longer discrete entities subject to isolated evaluation; they are deeply embedded within complex environments where computational processes, data flows, and physical systems interact continuously and dynamically. This integration is evident across multiple sectors. For example, in healthcare, a diagnostic algorithm operates within the same infrastructure that manages patient data, medical devices, and clinical protocols. Such interconnectedness presents a fundamental challenge for existing regulatory structures. Specifically, AI-related security failures often travel through both cybersecurity and AI-governance pathways, even when they can be described sequentially as a cyber incident followed by an AI system failure. For example, when an AI model is compromised through data poisoning (an attack in which malicious actors deliberately introduce

corrupted or manipulated data into the training dataset), the initial vulnerability may concern data collection, validation, or access control, while the downstream consequences concern model reliability, explainability, and accountability. It shows that AI has become the next technology-specific adaptation point for cybersecurity governance and that existing AI and cybersecurity frameworks need clearer coordination at the points where their duties interact. This growing complexity intensifies as agentic AI systems develop. Unlike conventional models, agentic systems can decompose goals into actions, use external application programming interfaces (APIs), maintain state, and adapt to feedback. This expanded operational capacity allows them to act across multiple systems and regulatory domains, making responsibility and compliance harder to assign.

In this context, the central problem is how regulatory frameworks built around separate governance assumptions can coordinate when AI-specific security challenges cut across their respective vocabularies, duties, and supervisory logics. The paper, therefore, uses “AI–cybersecurity alignment” to describe situations in which AI governance and cybersecurity governance increasingly address the same operational systems, evidence bases, and accountability problems. To analyze these developments, the paper develops the AI–Cybersecurity Taxonomy, a framework for identifying recurring governance functions that appear across otherwise distinct regulatory regimes.

Main Contributions of this Study

The paper makes three main contributions. First, it makes a taxonomic and methodological contribution by developing the AI–Cybersecurity Taxonomy through qualitative content analysis, comparative regulatory analysis, and grounded-theory-informed coding techniques, showing how differently framed instruments align around five recurring governance functions. Second, it makes a conceptual contribution by identifying the broader logics of provisional governance, agentic amplification, and embedded verifiability that explain both emerging alignment and continuing fragmentation. Third, it makes a practical and sectoral contribution by showing how the taxonomy helps interpret implementation challenges in critical infrastructure settings, especially in healthcare, energy, finance, and maritime operations.

The remainder of the paper develops these contributions as follows. The Literature Review outlines the theoretical foundations of technology regulation and the specific gap in AI-related cybersecurity governance; the Methodology section presents the qualitative content analysis, comparative regulatory analysis, and grounded-theory-informed coding approach; the Analysis section examines the five clusters; the Findings and Discussion section identifies the key patterns and their sectoral implications; the Implications for Research and Practice section considers the broader significance of the taxonomy; the Limitations and Future Directions section identifies the study’s constraints and areas for further research; the Recommendations section offers practical guidance for policymakers, regulators, standards bodies, and organizations; and the Conclusion closes the paper.

Literature Review

The governance literature consistently identifies coordination, resilience, and regulation-by-design as central challenges for emerging technologies (Brownsword & Yeung, 2008; Batool et al., 2025). A central way of framing that difficulty is to distinguish between design-based regulation, which embeds compliance and control mechanisms within technological architecture, and more conventional rule-based approaches that impose external behavioral mandates on developers and operators (Brownsword & Yeung, 2008). This distinction is especially relevant to AI-enabled systems and cybersecurity, where principles such as “security by design” and “security by default” reflect a broader shift toward integrating regulatory objectives into system design (Brownsword & Yeung, 2008; Yeung, 2017). Related studies on algorithmic regulation similarly treat technological systems not only as objects of regulation but also as instruments through which regulation is enacted, for example, through monitoring, constraint, and verification mechanisms embedded in technical environments (Yeung,

2018). Because emerging technologies are governed by multiple regulatory actors, instruments, and governance mechanisms, coordination has been proposed to link fragmented responses and address gaps without relying on a single comprehensive regime (Marchant & Wallach, 2015). Within that broader concern, cybersecurity scholarship has shown that governance involves not only preventing cyber incidents, but also coordinating actors, building resilience, and shaping how cyber threats are understood and addressed (Shackelford, 2012; Dunn Cavelty, 2014). Recent AI-governance reviews and comparative studies similarly show rapid proliferation of AI frameworks, risk-management models, and compliance tools, but they tend to focus on responsible AI governance broadly rather than on how cybersecurity regulation is adapting to AI-specific security challenges (Batool et al., 2025; Lobel, 2025; Finch & Butt, 2025). Work on AI in critical infrastructure shows that these systems generate operational and institutional risks when AI capabilities are embedded in infrastructure that also carries safety, security, and continuity obligations (Gerstein & Leidy, 2024). These strands of scholarship suggest that governance models come under strain when systems operate across interconnected infrastructure settings, adapt over time, and produce effects that are difficult to isolate within a single regulatory domain. In such conditions, governance cannot rely only on static classification or one-time compliance; it increasingly depends on ongoing monitoring, clearer human oversight, and accountability arrangements that can track how risks evolve in practice. Therefore, while the literature establishes the need for coordination, it leaves unresolved how that coordination is being operationalized. The specific gap addressed here concerns the regulatory treatment of AI-related cybersecurity problems. Existing scholarship has examined AI governance, cybersecurity resilience, and regulatory fragmentation, but it has not offered systematic analysis of how recent regulatory instruments translate AI-specific security concerns into recurring cybersecurity governance functions across jurisdictions and sectors. This study addresses that gap by comparing regulatory texts directly, using a taxonomy to identify where current frameworks align, where they adapt existing cybersecurity logics to AI-enabled systems, and where coordination remains incomplete. In particular, the present study approaches the field through a taxonomic lens. This analytical choice is informed by broader scholarship showing that regulation is often distributed, conceptually heterogeneous, and expressed through multiple overlapping forms rather than through a single coherent model (Black, 2002). It is also consistent with concept-centric synthesis in information systems research, which emphasizes organizing prior work around concepts and relationships rather than simply listing sources or authors (Webster & Watson, 2002). Finally, it aligns with work on regulation by design, which shows how governance increasingly operates through technical architecture, data environments, and behavioral shaping rather than through formal rules alone (Yeung, 2017).

Methodology

This study uses qualitative content analysis and comparative regulatory analysis to develop a taxonomy of AI-related cybersecurity governance functions. The coding process is grounded-theory-informed, by drawing on grounded-theory-informed techniques of iterative comparison and category refinement (Corbin & Strauss, 2015). Qualitative content analysis supports text-based interpretation and comparison across cases (Schreier, 2012), while comparative regulatory analysis identifies patterned similarities and differences across legal and policy instruments without assuming jurisdictional uniformity (Ragin, 1987). An iteratively updated codebook was used to refine category boundaries, address ambiguous cases, and maintain consistency across instruments. The analysis had three phases:

Phase 1: Regulatory Instrument Selection Criteria. Regulatory instruments were selected using four criteria: they had to address both AI and cybersecurity, date from 2022–2026 in terms of adoption, issuance, or publication, originate from authorities with significant jurisdictional reach or normative influence, and impose legal obligations or serve as closely relevant official standards or guidance. This range captured both foundational and recent developments shaping the current regulatory landscape. See Appendix A, Table A1 for the diverse sample, with balanced geographic and sectoral coverage.

Phase 2: Content Analysis and Multi-Dimensional Coding. The analysis coded regulatory provisions by function, governance objective, compliance mechanism, references to autonomous or agentic systems, and alignment with standard risk-management cycles. Primary sources included legislative texts, implementing regulations, technical standards, and official guidance. This source base captured both formal requirements and practical interpretive expectations.

Phase 3: Cluster Synthesis Using Grounded-Theory-Informed Coding Techniques (Open, Axial, Selective). The coded material was synthesized across all instruments to develop a taxonomy that shows how regulatory language aligns around recurring governance functions. Open coding first identified explicit regulatory terms and obligations. Axial coding then compared those codes across instruments to determine whether they addressed similar governance problems, such as risk stratification, lifecycle control, verifiability, incident response, or dependency oversight. Selective coding was used in a limited, taxonomic sense to consolidate those recurring patterns into the five clusters that organize the analysis. Although the Phase 2 coding dimensions necessarily shaped the analytic field, the clusters were not predetermined labels; they were retained only where multiple instruments, jurisdictions, and compliance mechanisms repeatedly pointed to the same governance function. The codebook was iteratively updated to record definitions, inclusion criteria, exclusion criteria, and illustrative examples to improve transparency on how the five clusters were derived (see Appendix B, Table B1).

Analysis

The analysis of the regulatory instruments resulted in the AI–Cybersecurity Taxonomy presented in Figure 1.



Figure 1: AI–Cybersecurity Taxonomy

The taxonomy is organized hierarchically around five core clusters: (1) Risk Classification and Dynamic Assessment, (2) Secure Lifecycle and Technical Controls, (3) Transparency, Explainability, and Verifiability, (4) Incident Response and Resilience, and (5) Supply Chain and Third-Party Stewardship. These clusters recur across major instruments from the European Union, the United States, and the Asia-Pacific region. Each cluster illustrates the harmonization of terminology, the industry concern it addresses, the regulatory sources in which the terms emerge, and the resulting risk-management logic.

Cluster 1: Risk Classification and Dynamic Assessment. Open coding identified recurring terms, such as “high-risk systems,” “high-impact systems,” “reasonable care,” and “risk-based proportionality,” each of which marks an effort to scale obligations to the level of harm a system may generate. For example, the EU Artificial Intelligence Act classifies systems as high-risk when they fall within specified product-safety or Annex III use cases and links that classification to heightened safeguards, while the Colorado Artificial Intelligence Act uses a reasonable-care standard supported by risk-management and impact-assessment duties, and California’s automated decision-making technology (ADMT) framework imposes additional duties when automated decision-making is used for significant decisions. Axial coding brought these differently framed provisions into relation by showing that they sort systems by expected harm and attach more demanding assessment, documentation, or oversight requirements as risk increases. At the same time, comparisons across instruments revealed the core governance challenge of fragmented classification, as thresholds for what counts as high-risk or high-impact vary, and cybersecurity considerations such as robustness, integrity, and resilience are incorporated unevenly. At the selective-coding stage, these recurring patterns were distilled into a broader logic of context-sensitive, dynamic risk proportionality, where classification is treated not as a fixed label but as an adaptive process requiring ongoing reassessment as systems evolve, operate in new environments, or encounter changing threat conditions. Thus, the cluster shows how risk classification functions as an integrative bridge between AI governance and cybersecurity risk management.

Cluster 2: Secure Lifecycle and Technical Controls. Open coding grouped terms, such as “security by design and by default,” “lifecycle technical controls,” “continuous monitoring,” and related references to end-to-end system integrity. These terms recur in instruments that treat security as something embedded across the life of the system rather than added after deployment. For example, the EU Cyber Resilience Act requires manufacturers of products with digital elements to ensure that products are “designed, developed and produced” in accordance with essential cybersecurity requirements, including security by design and by default, while the U.S. Food and Drug Administration (FDA) Medical Device Cybersecurity Guidance, the National Institute of Standards and Technology (NIST) Cybersecurity Framework 2.0, and the California Department of Technology SIMM 5305-A standard similarly emphasize ongoing security controls across development, deployment, and maintenance (California Department of Technology, 2025). Axial coding showed that, despite sectoral and jurisdictional differences, these provisions converge on the same governance challenge: AI systems remain exposed to adversarial threats throughout design, training, deployment, operation, and decommissioning, and those risks intensify when systems continue learning, updating, or acting autonomously after deployment. Selective coding framed these features as a logic of continuous lifecycle security, in which protection is embedded across the full system lifecycle and maintained through ongoing, system-wide controls rather than one-time certification or static safeguards.

Cluster 3: Transparency, Explainability, and Verifiability. Open coding grouped terms such as “technical documentation,” “plain-language explanation,” “human review,” “auditability,” and “stakeholder disclosure,” which together signal a regulatory demand for interpretable and verifiable AI behavior. These terms appear in different legal forms but serve related functions. For example, the EU Artificial Intelligence Act requires providers of high-risk AI systems to prepare technical documentation and maintain records sufficient to demonstrate compliance, while the California Consumer Privacy Act’s ADMT rules require notice, access, and explanation rights in relation to significant automated decisions, and the New York RAISE Act introduces transparency and safety-related obligations for

covered frontier AI systems. Axial coding emphasized a shared governance problem across these provisions: algorithmic opacity limits not only understanding of how systems produce outputs, but also the ability to assess integrity, detect manipulation, trace anomalies, and assign responsibility in adversarial environments. Selective coding reorganized these requirements as a logic of verifiable transparency, in which explainability and auditability become enforceable structural requirements that support both meaningful oversight of AI decision-making and the capacity to monitor, test, and secure systems against manipulation or failure.

Cluster 4: Incident Response and Resilience. Open coding grouped terms such as “incident notification,” “major information and communication technology (ICT)-related incident,” “root-cause analysis,” “cyber resilience,” and related reporting and recovery obligations. These terms indicate that incidents are being regulated not simply as isolated failures but as events that reveal system vulnerabilities and governance weaknesses. For example, the Digital Operational Resilience Act (DORA) requires financial entities to provide initial notification within four hours after classifying an ICT-related incident as major, and no later than twenty-four hours after detection or awareness, while the Network and Information Security Directive 2 (NIS2) requires notification within twenty-four hours. The U.S. Coast Guard Maritime Cybersecurity Rule mandates immediate reporting when a cyber event affects the safety or security of a vessel or the marine environment. Axial coding showed that, across instruments, these provisions address a common governance challenge: in distributed, partially autonomous environments, incidents may stem from malicious interference, system design limitations, data anomalies, or autonomous behavior, making detection, attribution, and timely response more difficult. Selective coding brought these features together as structured resilience governance, in which reporting, investigation, coordination, and system adaptation are formalized as connected obligations designed to strengthen robustness over time rather than simply restore operations after failure.

Cluster 5: Supply Chain and Third-Party Stewardship. Open coding grouped terms such as “supply chain and third-party risk management,” “Software Bill of Materials (SBOM),” “component provenance tracking,” “coordinated vulnerability disclosure,” and related obligations concerning external providers and dependencies. These terms repeatedly appeared where regulators sought to extend visibility and accountability beyond a single organization. For example, the EU Cyber Resilience Act requires a Software Bill of Materials and coordinated vulnerability disclosure, DORA formalizes third-party risk management through contractual oversight and concentration-risk assessment, and the FDA guidance requires software-component documentation and SBOM-related disclosure for covered medical devices, including devices that incorporate artificial intelligence and machine learning (AI/ML) functionality. Axial coding highlighted that these provisions respond to the same governance challenge: AI systems depend on complex, multi-tier networks of software, models, datasets, and service providers, but regulators and organizations often lack sufficient visibility into where vulnerabilities originate and how they propagate across those dependencies. At the selective-coding level, this pattern was synthesized into a regulatory logic of ecosystem-wide accountability, in which enforceable information flows, traceability requirements, and third-party governance mechanisms are used to strengthen both cybersecurity resilience and trust in AI-enabled systems.

Findings and Discussion

As shown in Figure 1 and the cluster analysis, the paper identifies several areas where AI and cybersecurity regulation are already aligned. Across the five clusters, different instruments increasingly converge on the need to classify systems by risk, embed security across the lifecycle, make system behavior more transparent and verifiable, formalize incident response, and strengthen oversight of external dependencies. These alignments are important because they show that regulators in different jurisdictions are beginning to respond to a shared practical reality: AI-enabled systems in critical infrastructure cannot be governed effectively through separate AI and cybersecurity logics. At the same time, the comparison also reveals that this alignment remains partial. Similar governance functions are

expressed through different legal vocabularies, triggered by different thresholds, and assigned to different institutional pathways. For example, Cluster 1 shows that the EU Artificial Intelligence Act, the Colorado Artificial Intelligence Act, and California's ADMT framework all scale obligations to harm, but they do so through different categories, triggers, and compliance duties. Cluster 4 shows a similar pattern in incident governance: DORA, NIS2, and the U.S. Coast Guard Maritime Cybersecurity Rule all require rapid reporting, but they define triggering events and timelines differently. Cluster 5 similarly reveals convergence around dependency oversight, but through distinct mechanisms such as software bills of materials, third-party risk management, and software-component documentation. What remains underdeveloped are the connecting mechanisms: shared classification criteria, coordinated triggers, interoperable reporting, and clearer rules for reassessment and intervention. The central finding is not simply convergence, but convergent problem recognition without integrated regulatory design.

That partial alignment becomes most visible at the point of implementation. Because the relevant duties remain distributed across separate instruments and supervisory traditions, organizations still face fragmentation when they try to operationalize compliance. A system may be classified as high-risk under one instrument, subject to security-by-design obligations under another, and required to report incidents under a third, without any explicit mechanism linking those duties into a single operational framework. Seen in this light, the analysis clarifies not only where regulation is beginning to align, but also why implementation remains difficult even when substantive concerns increasingly overlap.

The taxonomy points to two substantive shifts in regulatory logic. The first is what this study terms *provisional governance*, that is, governance designed on the assumption that systems, risks, and compliance conditions will change over time. This pattern appears most clearly in Cluster 1 through dynamic risk assessment, in Cluster 2 through lifecycle-based controls and continuous monitoring, in Cluster 4 through incident obligations that require investigation and adaptation after initial reporting, and in Cluster 5 through traceability duties that must remain current as dependencies change. The second is the *agentic amplification effect*, by which agentic capabilities intensify familiar risks by extending their scale, speed, and organizational reach. Reflected across the clusters, features associated with agentic systems (task decomposition, tool use, memory, and adaptive behavior) do not simply create new risks; they make existing vulnerabilities harder to isolate, contain, and govern once multiple functions, infrastructures, and decision layers become involved. This dynamic helps explain why the sectoral examples in healthcare, energy, finance, and maritime settings are significant: a vulnerability that might once have remained local can now produce wider operational consequences, making the combined governance logics identified by the taxonomy consequential in practice.

These shifts also help explain the clearest unresolved gap in the current framework: control. As anticipated in the Literature Review and confirmed by the taxonomy, current instruments increasingly recognize autonomy of agentic capabilities but rarely specify how governance should operate once a system can modify its behavior in ways that are difficult to anticipate *ex ante*. This gap is visible across multiple clusters simultaneously. Cluster 1 shows that risk classification is becoming more dynamic, but most instruments still do not specify when a changing system should be reclassified. Cluster 3 shows a growing emphasis on explainability and auditability, but these mechanisms do not, by themselves, determine when intervention is required. Cluster 4 formalizes reporting and investigation duties, but still leaves uncertainty about how responsibility should be attributed when harmful behavior emerges gradually rather than through a discrete failure. A related difficulty is the *learning-authority dilemma*: once adaptive systems continue to learn, update, or act across distributed settings, governance may recognize the need for oversight without clearly specifying how authority should be retained, exercised, or reasserted as behavior changes. This dilemma exposes a structural weakness running across the taxonomy: many instruments acknowledge adaptive systems but few articulate governance models capable of retaining authority once behavior becomes dynamic, distributed, and partially self-directed.

These findings point to an emerging regulatory rationality centered on embedded verifiability, which envisages governance arrangements that integrate monitoring, traceability, and intervention into AI

systems, rather than treating verification as a purely external or documentary exercise. The significance of that rationality becomes clearest when the taxonomy is read as a whole rather than as five isolated themes. Figure 1 shows that the clusters form an integrated governance structure in which classification, lifecycle controls, transparency, incident response, and supply-chain stewardship each create different points at which AI-enabled systems can be monitored, tested, and, where necessary, interrupted. The taxonomy's contribution is not only to map convergence, but to show how current instruments are moving toward integrated governance while still lacking some of the legal and institutional mechanisms needed. The sectoral considerations that follow develop this point further by showing how the same alignments, implementation strains, and control gaps manifest differently across different settings.

Sectoral Considerations

The sectoral examples clarify the practical significance of the taxonomy's findings. As the previous section shows, the taxonomy identifies concrete points at which AI and cybersecurity regulation are beginning to align, while showing that these alignments remain incomplete. Sectoral analysis makes that tension visible in operational terms by showing where the five clusters map onto real-world infrastructure, where provisional governance and agentic amplification become most prominent, and where unresolved questions of coordination, timing, and responsibility persist. These examples are illustrative, not empirical case studies, and serve to contextualize the regulatory patterns identified.

Healthcare and energy show this pattern under different operating conditions. In healthcare, Cluster 1 focuses on risk-based treatment of high-impact clinical systems, Cluster 2 on secure lifecycle expectations for connected devices and data integrity, and Cluster 3 on explainability and auditability where decisions affect patient safety. The EU Artificial Intelligence Act's high-risk and technical documentation duties therefore directly intersect with the FDA's cybersecurity expectations for AI/ML-enabled medical devices. However, legacy infrastructure and procurement dependencies complicate implementation. In operational terms, a compromised imaging dataset or an insecure interface between a diagnostic model and an electronic health record may require clinical review, cyber containment, documentation updates, and vendor coordination at the same time. In energy, the same tension arises under extreme temporal pressure: grid optimization, predictive maintenance, and renewable balancing bring Cluster 1 and Cluster 4 into particularly close relation, because risk classification must remain dynamic while incident governance must operate faster than conventional reporting structures were designed to allow. An AI-driven grid reconfiguration exploit, for example, may manipulate sensor inputs or optimization parameters so that the system reroutes load during a narrow operating window before routine human review can occur. NIS2 and the EU Cyber Resilience Act provide important alignment points through incident reporting, resilience, and security-by-design obligations, but, as the broader discussion noted, they still do not resolve how oversight should operate when AI-driven reconfiguration occurs faster than human review or ordinary supervisory timelines permit.

Financial services and maritime operations further show partial alignment, combined with implementation strain. In finance, Clusters 3 and 5 are especially important because transparency and verifiability support both market accountability and system assurance, while third-party risk management becomes essential in environments that depend on external models, cloud services, and payment infrastructures. DORA creates one of the clearest frameworks in the dataset by linking resilience, incident reporting, and third-party oversight within a single regime; even so, the speed of automated financial activity compresses the time available for review, attribution, and intervention, so formal compliance windows do not always map cleanly onto operational reality. Maritime operations reveal a similar gap in the interaction between new AI capabilities and older safety architectures. Clusters 2 and 4 are especially salient: secure lifecycle controls are weakened by legacy infrastructure, and incident response must address both cyber disruption and physical safety consequences. The U.S. Coast Guard Maritime Cybersecurity Rule begins to bridge these domains, but safety certification, cybersecurity assurance, and AI-specific oversight often remain layered rather than integrated.

Across all four sectors, then, the pattern anticipated in the broader discussion holds: the closer AI systems move to core operational functions, the less sustainable it becomes to regulate AI governance and cybersecurity as separate domains. Cross-domain instruments such as DORA, the EU Cyber Resilience Act, the FDA Medical Device Cybersecurity Guidance, and the U.S. Coast Guard Maritime

Cybersecurity Rule illustrate movement toward combined governance. However, they also confirm that misaligned thresholds, incomplete coordination, and uncertainty about intervention and responsibility continue to reappear in practice.

Implications for Research and Practice

The taxonomy's theoretical significance is best understood in relation to the frameworks discussed in the Literature Review. First, it extends the distinction between rule-based and design-based regulation by showing that governance increasingly depends on how regulatory objectives are embedded in system design, monitoring, and operational controls rather than imposed only as external legal commands (Brownsword & Yeung, 2008; Yeung, 2017). Second, it gives empirical substance to scholarship on algorithmic regulation by showing how contemporary instruments distribute regulatory functions across classification, lifecycle security, verifiability, resilience, and dependency management (Yeung, 2018). Third, it develops Marchant and Wallach's account of governance coordination by identifying the concrete points at which coordination is already emerging across domains, as well as the institutional gaps that still prevent those points of alignment from becoming coherent governance (Marchant & Wallach, 2015). Finally, the taxonomy also speaks to broader work on cyber resilience and regulatory fragmentation by showing that current legal systems are increasingly being reorganized around operational and institutional infrastructure rather than around previously separate categories of AI risk and cyber risk (Shackelford, 2012; Dunn Caveltly, 2014; Black, 2002). For practice, the taxonomy's value lies less in prescribing a single regulatory model than in helping actors work across the implementation problems identified in the previous discussion. For regulators, it offers a structured way to see where existing duties already align across the five clusters and where they still break down because of misaligned thresholds, fragmented supervisory pathways, or uncertainty about intervention and responsibility. For regulated entities, especially in critical infrastructure sectors, it provides a basis for mapping compliance obligations that currently sit across separate AI and cybersecurity regimes and for identifying where integrated controls are needed in areas such as risk reassessment, incident escalation, traceability, and third-party oversight. For researchers, it provides a comparative vocabulary for studying how these issues vary across sectors and jurisdictions.

Limitations and Future Directions

This study presents several limitations that also point to opportunities for further development. First, the sample was deliberately bounded to regulatory instruments adopted, issued, or published between 2022 and 2026, with an emphasis on major jurisdictions and selected instrument types. This design made close comparison of recent and influential developments possible and supported the development of a coherent taxonomy. Future research could extend this foundation by adding other national, subnational, sector-specific, and voluntary materials. Second, the study uses qualitative content analysis, comparative regulatory analysis, and grounded-theory-informed coding techniques to derive the five clusters inductively from regulatory language. This method was well suited to identifying emerging patterns in a rapidly evolving field and to developing an empirically grounded taxonomy. At the same time, like other interpretive approaches, it offers one analytically coherent way of organizing the material. Alternative methods could therefore complement the present framework, test the robustness of the clusters, or refine their boundaries. Third, the analysis concentrates on formal legal texts, technical standards, and official guidance. This textual focus was especially valuable for examining regulatory design, intent, and conceptual structure, while leaving room for future work on how the identified points of alignment and fragmentation develop in implementation. A useful next step would be to examine these patterns in organizational settings. Future researchers could use the codebook as a basis for interviews with Chief Information Security Officers (CISOs), AI governance leads, compliance officers, and sector regulators to test whether the five clusters correspond to actual

budgeting decisions, organizational restructuring, reporting practices, and control ownership. Case studies, compliance audits, enforcement analysis, and stakeholder interviews would therefore offer useful extensions of the taxonomy. Finally, the regulatory landscape addressed in this paper remains highly dynamic, particularly in relation to agentic and multimodal systems. The taxonomy should thus be read as a baseline account of the regulatory environment. As new regulatory instruments emerge and existing frameworks are interpreted through practice, continued application, revision, and comparative testing of the taxonomy will help keep the framework responsive to ongoing developments.

Recommendations

The taxonomy suggests that the main regulatory task is no longer to create separate bodies of law for AI and cybersecurity, but to make existing and emerging regimes work together around the same operational realities. The recommendations below translate the five clusters into practical directions for policymakers, regulators, standards bodies, and regulated entities. They are designed to be modular rather than uniform: shared in structure across sectors, but adaptable in implementation to different infrastructures, supervisory models, and risk profiles.

1. Build interoperable governance architecture around shared functions. The taxonomy shows that risk classification, technical security, transparency, incident response, and dependency management already recur across multiple instruments, even when expressed through different legal vocabularies. One practical implication is that regulators and standards bodies should use those shared functions as the organizing logic for coordination by aligning core definitions such as system integrity, material incident, harmful system behavior, and third-party dependency across AI and cybersecurity regimes, issuing joint interpretive guidance where obligations overlap, and designing compliance templates that allow the same evidence base to satisfy related duties across frameworks. At the organizational level, the same logic can be reflected by integrating AI governance and cybersecurity ownership into shared control frameworks, common risk registers, and cross-functional review processes rather than treating them as separate compliance streams. In healthcare, for example, hospitals can align clinical AI oversight with medical-device cybersecurity review, while in finance, firms can connect AI model governance with operational resilience and incident management teams. This recommendation follows directly from Cluster 1 and Cluster 4, where partial alignment is already visible but still breaks down because thresholds, triggers, and reporting pathways remain fragmented.

2. Shift from static compliance to adaptive supervision and adaptive operations. A central implication of the taxonomy is that governance must track systems that change over time. This calls for regulators and supervisors to supplement one-time certification and periodic reporting with mechanisms that trigger reassessment when systems are retrained, materially updated, connected to new data sources, or deployed in new operational environments. Parallel internal practices are equally important: organizations should define change thresholds that trigger renewed review, document model and system updates in ways that support both technical assurance and compliance, and link deployment decisions to continuous monitoring rather than relying solely on fixed approval points. In energy, this may mean reassessing AI-enabled grid optimization tools when operating conditions or data feeds change, while in healthcare it may involve renewed review when clinical decision-support systems are retrained or integrated into new diagnostic workflows. This recommendation reflects the logic of provisional governance identified in the Findings and Discussion. It is especially important for Cluster 1 and Cluster 2, where the analysis showed that risk classification and lifecycle controls increasingly depend on continuous recalibration rather than fixed compliance moments.

3. Govern agentic systems through intervention points, not only design constraints. The learning-authority dilemma identified in this study suggests that oversight cannot depend solely on ex-ante restrictions where systems can initiate sequences of actions across multiple tools or environments. Regulators should require that high-impact agentic systems include enforceable intervention points:

predefined conditions under which actions can be paused, escalated for human review, or automatically contained. Organizations should operationalize this expectation by defining escalation rules, kill-switch or containment procedures, authority boundaries for autonomous actions, and review mechanisms for external tool use and delegated tasks. In maritime operations, for example, this operational need may include manual override points for AI-assisted navigation or collision-avoidance systems, while in financial services, it may involve escalation thresholds for autonomous trading or fraud-response tools that interact with payment infrastructure. The aim is not to eliminate autonomy, but to preserve governability once autonomous action becomes operationally significant. This recommendation draws most directly on Clusters 1, 3, and 4, where the analysis revealed persistent uncertainty about reclassification, intervention, and responsibility when behavior changes over time.

4. Move from documentary transparency to operational verifiability. The taxonomy shows that transparency is increasingly expected to serve not only accountability goals, but also security and resilience functions. Regulators and standards bodies should therefore place greater emphasis on auditability that is machine-executable, testable, and linked to system operations, rather than treating it only as a documentation requirement. Organizations should support this shift by implementing logging for high-impact decisions, traceable records of model updates and external tool use, and reviewable links between model behavior and cyber telemetry so that transparency becomes useful for assurance and response rather than purely formal. In healthcare, this operational capability can support post-event review of AI-assisted diagnostic decisions. In finance, it can help connect trading or fraud-detection outputs to the system events and external calls that produced them. Cluster 3 provides the clearest basis for this recommendation, while Cluster 4 shows why such capabilities are necessary for timely investigation and response when incidents involve both algorithmic and security-related failures.

5. Extend dependency governance beyond software components. Cluster 5 shows that AI-related risk increasingly spreads across models, datasets, external APIs, cloud infrastructure, and service providers, rather than solely through software. For that reason, regulators should build on existing software bill-of-materials approaches by encouraging broader forms of dependency mapping that capture model provenance, material data dependencies, external services, and critical update pathways. Organizations can reinforce the same objective by maintaining inventories of critical models, datasets, vendors, and external integrations, and by linking procurement, vendor review, and incident planning to those dependency maps. In maritime settings, for example, this may include tracking dependencies on navigation data services and legacy communications vendors, while in healthcare it may involve mapping third-party datasets, cloud-hosted clinical AI tools, and connected device suppliers. The objective should be proportionate traceability: enough visibility to identify where vulnerability, bias, or failure may propagate, without assuming that full transparency is always possible in complex ecosystems. This recommendation is especially important in sectors such as healthcare, finance, and maritime operations, where the sectoral analysis showed that legacy systems and third-party dependencies often determine how failures spread in practice.

6. Combine shared principles with sector-specific implementation and organizational capability. This recommendation is supported by the combined findings of all five clusters, each of which highlights a distinct dimension of regulatory adaptation across sectors. Cluster 1 demonstrates that risk thresholds must be adapted to context; Cluster 2 highlights that lifecycle controls need to reflect sector-specific systems and update cycles; Cluster 3 shows that verifiability depends on the evidence available within a given operational setting; Cluster 4 indicates that incident response mechanisms must align with sectoral speed and safety consequences; and Cluster 5 illustrates that dependency governance varies according to supply-chain structure. Therefore, regulators should preserve common dual-domain principles across sectors, while allowing implementation to vary through sector-specific technical standards, supervisory guidance, certification pathways, and incident protocols. Organizations should translate those principles into sector-appropriate governance arrangements, including cross-functional response teams, tailored assurance processes, and clear escalation pathways that reflect their operational environment. For example, an energy operator may need real-time coordination among control-room

engineers, cybersecurity teams, and AI vendors, while a hospital may need escalation pathways linking clinicians, procurement, data protection, and device security teams. International standards bodies, including ISO/IEC, can serve as a neutral translation layer by expressing these shared functions in interoperable control language that different jurisdictions and sectors can adapt without requiring identical legal triggers (International Organization for Standardization & International Electrotechnical Commission, 2023a, 2023b). At the same time, capacity support remains essential, especially for smaller operators and public-sector organizations. Shared sandboxes, reference control frameworks, model documentation templates, pooled audit resources, and industry coordination mechanisms can help make dual-domain compliance more realistic without lowering substantive expectations. These recommendations suggest that the next step in AI–cybersecurity governance is not an extensive replacement of existing frameworks, but a more deliberate effort by both regulators and organizations to connect them around the five functional areas identified by the taxonomy. Therefore, the goal is to make governance more interoperable, adaptive, and intervention-capable as AI-enabled systems become more embedded, autonomous, and infrastructural.

Conclusion

This paper has shown that the conventional separation between AI regulation and cybersecurity regulation no longer reflects how contemporary AI-enabled systems are actually governed in practice. The AI–Cybersecurity Taxonomy makes this change visible by identifying the points at which these domains now intersect, align, and still remain only partially coordinated. Across the instruments examined, regulatory attention increasingly aligns around recurring governance functions captured by the five clusters. They show that AI is becoming a technology-specific adaptation point for cybersecurity governance and that regulatory attention is increasingly organized around adaptive, interdependent, and infrastructural systems. At the same time, the taxonomy clarifies why alignment remains incomplete. Regulatory frameworks still rely on categories developed for more stable, separable systems than those now being governed, so duties that are increasingly aligned in purpose remain fragmented in classification, supervision, and intervention. One consequence is that governance may appear comprehensive on paper while remaining discontinuous in practice, a problem that becomes especially evident when autonomous action, external dependencies, and sector-specific infrastructures interact. The broader implication is that AI and cybersecurity governance are no longer best understood as parallel tracks, but as interdependent dimensions of the same regulatory space. If so, the question for future governance is not only how to coordinate rules, but also how to preserve institutional attention in environments where systems evolve faster than categories, and where failure may emerge gradually before becoming visible as a reportable event. In that respect, the taxonomy is useful not only because it maps where existing frameworks already coordinate, but also where they remain fragmented and where regulatory blind spots may form next.

Acknowledgement

This research was supported, in part, by the Cybersecurity at MIT Sloan (CAMS) initiative (<https://cams.mit.edu>).

References

- Australian Government. (2025). *Voluntary AI safety standard*.
<https://www.industry.gov.au/publications/voluntary-ai-safety-standard>
- Batool, A., Zowghi, D., & Bano, M. (2025). AI governance: A systematic literature review. *AI and Ethics*, 5(3), 3265–3279. <https://doi.org/10.1007/s43681-024-00653-w>

- Black, J. (2002). Critical reflections on regulation. *Australian Journal of Legal Philosophy*, 27, 1–35. <https://www.austlii.edu.au/au/journals/AUJILawPhil/2002/1.pdf>
- Brownsword, R., & Yeung, K. (2008). *Regulating technologies: Legal futures, regulatory frames and technological fixes*. Hart Publishing. <https://www.bloomsbury.com/us/regulating-technologies-9781841137889/>
- California Department of Technology. (2025). *SIMM 5305-A Information Security Program Management Standard*. Statewide Information Management Manual. <https://www.cdt.ca.gov/wp-content/uploads/2025/08/SIMM-5305-A-Information-Security-Program-Management-Standard-20250918-FINAL.pdf>
- California Privacy Protection Agency. (2025). *California Consumer Privacy Act automated decision-making technology (ADMT) regulations*. https://cpa.ca.gov/regulations/ccpa_updates.html
- Cisco. (2026). *State of industrial AI report 2026*. <https://www.cisco.com/c/dam/en/us/solutions/networking/industrial-iot/2026-state-of-industrial-ai-report.pdf>
- Colorado General Assembly. (2024). *SB24-205: Consumer protections for artificial intelligence*. <https://leg.colorado.gov/bills/sb24-205>
- Corbin, J., & Strauss, A. (2015). *Basics of qualitative research: Techniques and procedures for developing grounded theory* (4th ed.). Sage. <https://collegepublishing.sagepub.com/products/basics-of-qualitative-research-4-235578>
- Digital Agency of Japan. (2025). *Act on the Promotion of Research and Development, and Utilization of AI-Related Technology (AI Act)*. <https://www.japaneselawtranslation.go.jp/en/laws/view/5066>
- Dunn Cavelty, M. (2014). *Cyber-security and threat politics: U.S. efforts to secure the information age*. Routledge. <https://www.routledge.com/Cyber-Security-and-Threat-Politics-US-Efforts-to-Secure-the-Information-Age/Dunn/p/book/9780415569880>
- European Union. (2022). *Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS2 Directive)*. Official Journal of the European Union. <https://eur-lex.europa.eu/eli/dir/2022/2555/oj>
- European Union. (2022). *Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA)*. Official Journal of the European Union. <https://eur-lex.europa.eu/eli/reg/2022/2554/oj>
- European Union. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- European Union. (2024). *Regulation (EU) 2024/2847 on horizontal cybersecurity requirements for products with digital elements (Cyber Resilience Act)*. Official Journal of the European Union. <https://eur-lex.europa.eu/eli/reg/2024/2847/oj>
- Gerstein, D. M., & Leidy, E. N. (2024). *Emerging technology and risk analysis: Artificial intelligence and critical infrastructure* (RR-A2873-1). RAND Corporation. https://www.rand.org/pubs/research_reports/RRA2873-1.html

- Finch, W. W., & Butt, M. (2025). Gaps in AI-compliant complementary governance frameworks' suitability (for low-capacity actors), and structural asymmetries (in the compliance ecosystem)—A systematic review. *Journal of Cybersecurity and Privacy*, 5(4), 101. <https://doi.org/10.3390/jcp5040101>
- International Organization for Standardization, & International Electrotechnical Commission. (2023a). *ISO/IEC 23894:2023 artificial intelligence — guidance on risk management*. <https://www.iso.org/standard/77304.html>
- International Organization for Standardization, & International Electrotechnical Commission. (2023b). *ISO/IEC 42001:2023 artificial intelligence management systems*. <https://www.iso.org/standard/81230.html>
- Lobel, O. (2025). The AI regulatory pyramid: A taxonomy and analysis of the emerging toolbox in the global race for the regulation and governance of artificial intelligence. *Loyola of Los Angeles Law Review*, 57(4), 859–914. <https://digitalcommons.lmu.edu/llr/vol57/iss4/1/>
- Marchant, G. E., & Wallach, W. (2015). Coordinating governance of emerging technologies. *Issues in Science and Technology*, 31(4), 43–48. <https://issues.org/coordinating-technology-governance/>
- Marotta, A., & Madnick, S. (2021). Convergence and divergence of regulatory compliance and cybersecurity. *Issues in Information Systems*, 22(1), 10–50. https://doi.org/10.48009/1_iis_2021_10-50
- Ministry of Science and ICT. (2025). *Basic Act on the Development of Artificial Intelligence and the Establishment of a Trustworthy Foundation*. <https://www.msit.go.kr/eng/bbs/view.do?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1214&sCode=eng>
- National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. <https://www.nist.gov/itl/ai-risk-management-framework>
- National Institute of Standards and Technology. (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1)*. <https://doi.org/10.6028/NIST.AI.600-1>
- National Institute of Standards and Technology. (2024). *Cybersecurity Framework (CSF) 2.0*. <https://www.nist.gov/cyberframework>
- New York State Senate. (2025). *Responsible AI Safety and Education (RAISE) Act*. <https://www.nysenate.gov/legislation/bills/2025/A6453>
- Ragin, C. C. (1987). *The comparative method: Moving beyond qualitative and quantitative strategies*. University of California Press. https://books.google.com/books/about/The_Comparative_Method.html?id=mZi17vherScC
- Schreier, M. (2012). *Qualitative content analysis in practice*. Sage. <https://uk.sagepub.com/en-gb/eur/qualitative-content-analysis-in-practice/book234633>
- Shackelford, S. J. (2012). *Managing cyber attacks in international law, business, and relations: In search of cyber peace*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139021838>

UK Government. (2024). *Implementing the UK's AI Regulatory Principles: Initial Guidance for Regulators*.

https://assets.publishing.service.gov.uk/media/65c0b6bd63a23d0013c821a0/implementing_the_uk_ai_regulatory_principles_guidance_for_regulators.pdf

U.S. Coast Guard. (2025). *Maritime cybersecurity rule*.

<https://www.federalregister.gov/documents/2025/01/17/2025-00708/cybersecurity-in-the-marine-transportation-system>

U.S. Food and Drug Administration. (2026). *Cybersecurity in medical devices: Quality system considerations and content of premarket submissions*.

<https://www.fda.gov/media/119933/download?attachment>

Webster, J., & Watson, R. T. (2002). Analyzing the past to prepare for the future: Writing a literature review. *MIS Quarterly*, 26(2), xiii–xxiii. <https://doi.org/10.2307/4132319>

Yeung, K. (2017). ‘Hypernudge’: Big data as a mode of regulation by design. *Information, Communication & Society*, 20(1), 118–136. <https://doi.org/10.1080/1369118X.2016.1186713>

Yeung, K. (2018). Algorithmic regulation: A critical interrogation. *Regulation & Governance*, 12(4), 505–523. <https://doi.org/10.1111/rego.12158>

Appendix A. Regulatory instrument sample

Appendix A presents the regulatory sample used in the study and how the instruments were compared across jurisdictions. Table A1 identifies each instrument’s jurisdiction, primary focus, key dates, and degree of explicit coverage of autonomous or agentic AI. The date and coverage fields clarify each instrument’s temporal status and whether autonomous systems are addressed directly or only through broader risk, security, or infrastructure provisions. The table supports the empirical basis, transparency, documentary validation, and later analysis of alignment and gaps in AI–cybersecurity governance.

Table A1. Regulatory Instrument Sample (ordered by the first relevant date for each instrument)

Regulation / Instrument	Jurisdiction	Primary Focus	Key Operational Dates	Autonomous/Agentic Coverage
U.S. Food and Drug Administration (FDA) Medical Device Cybersecurity Guidance	United States (Federal)	Cybersecurity for medical devices, including cyber devices and AI/ML-enabled devices	Final guidance issued in February 2026; supersedes the June 2025 version and updates the September 2023 guidance	Explicit (AI/ML in medical devices)
ISO/IEC 23894:2023 artificial intelligence — guidance on risk management	International	AI-specific risk management guidance relevant to AI-enabled systems	Published in February 2023	Implicit (through AI risk-management guidance)
National Institute of Standards and Technology (NIST) AI Risk Management Framework (AI RMF 1.0)	United States (Federal)	Voluntary AI governance framework incorporating safety, trustworthiness, and security	Published in January 2023; widely referenced in 2024–2026	Explicit (AI system governance)

Issues in Information Systems

Volume 27, Issue 2, pp. 419-436, 2026

ISO/IEC 42001:2023 AI Management System	International	Management system for AI with security controls	Published in 2023; widespread adoption in 2025–2026	Explicit (AI system management)
National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF) 2.0	United States (Federal)	Cybersecurity framework relevant to AI-enabled systems and digital infrastructure	Released in February 2024	Explicit (AI-specific security)
NIST AI 600-1: Generative Artificial Intelligence Profile	United States (Federal)	Generative AI risk management, including security, provenance, and incident disclosure	Published in July 2024	Explicit (generative AI systems)
UK AI Regulatory Principles Guidance	United Kingdom	Principles-based safety, security, and robustness guidance for regulators	Published in February 2024; implementation ongoing through sector regulators	Implicit (sector-specific)
EU Digital Operational Resilience Act (DORA)	European Union	ICT risk management and operational resilience for the financial sector	Fully applicable from 17 January 2025	Implicit (third-party AI services)
EU Network and Information Security Directive 2 (NIS2)	European Union	Cybersecurity for essential entities using AI	Member State transposition deadline: 17 October 2024; national implementation remains uneven	Implicit (critical infrastructure)
Australia Voluntary AI Safety Standard	Australia	Voluntary guardrails for safe AI deployment	Released in 2025; adoption continuing into 2026	Explicit (AI safety requirements)
California Cybersecurity Audit Regulations	California, USA	Cybersecurity audit obligations relevant to AI-enabled systems	Effective 1 January 2026; initial audit certifications due 2028–2030 depending on revenue	Implicit (where AI systems fall within audited environments)
Act on the Promotion of Research and Development, and Utilization of AI-Related Technology (AI Act)	Japan	Light-touch cooperation and safety measures	Enacted in May 2025	Implicit (government coordination)
New York RAISE Act	New York, USA	Transparency and safety for frontier AI	Signed in December 2025; implementation expected in 2027, subject to amendments	Explicit (frontier AI systems)
South Korea AI Framework Act (Basic Act on Development of AI and Establishment of a Trustworthy Foundation)	South Korea	Transparency and safety for AI systems	Passed in January 2025; implementation from 2026	Explicit (AI system governance)

U.S. Coast Guard Maritime Cybersecurity Rule	United States (Federal)	Cybersecurity for maritime transportation (including AI)	Effective from 16 July 2025	Implicit (maritime operations)
EU Artificial Intelligence Act (AI Act)	European Union	Risk-based AI governance, including robustness and cybersecurity obligations for high-risk systems	Progressive application from 2 August 2026, with later application for some high-risk categories	Implicit (via high-risk classification)
Singapore Model AI Governance Framework for Generative and Agentic AI	Singapore	Risk management integrating cybersecurity (agentic focus)	Launched in January 2026; updated framework version published in May/June 2026	Explicit (agentic systems)
California Consumer Privacy Act – automated decision-making technology (ADMT) regulations	California, USA	Privacy and security for automated decision-making	General updates effective 1 January 2026; ADMT requirements from 1 January 2027	Explicit (automated decision-making technology)
Colorado Artificial Intelligence Act	Colorado, USA	Risk management for high-risk AI systems	Effective from 1 February 2026	Explicit (high-risk AI systems)
EU Cyber Resilience Act (CRA)	European Union	Cybersecurity for products with digital elements (including AI-enabled products)	Reporting from 11 September 2026; full application from 11 December 2027	Implicit (product security requirements)

Appendix B. Codebook extract

Appendix B illustrates how the codebook was used in the analysis. As an example, Table B1 shows how recurring Cluster 1 terms were translated into comparable coding categories by defining what each term meant, what counted for inclusion or exclusion, and which regulatory examples supported the coding decisions. The extract is included to strengthen the internal transparency of the taxonomy: it shows how textual provisions were operationalized as codes, how coding boundaries were drawn, and how examples from the regulatory corpus supported the transition from individual terms to broader cluster-level governance functions. Although the full coding process involved all five clusters, this extract provides a replicable template for assessing the logic used across the analysis.

Table B1. Codebook Extract (Cluster 1 Example)

Code Category	Definition	Inclusion Criteria	Exclusion Criteria	Illustrative Regulatory Excerpts
Risk-Based Proportionality Logic	Principle linking regulatory obligations to level of risk	Any formulation of scaling obligations based on harm level	Compliance/no-compliance regime	EU AI Act: obligations and safeguards calibrated to the level of risk of the system
High-Risk Classification	Formal legal designation of AI systems based on potential harm	Any explicit “high-risk,” “high-impact,” or equivalent classification	General AI references without risk stratification	EU AI Act Art. 6: high-risk classification applies to specified product-safety and Annex III use cases, with

		triggering regulatory obligations		associated heightened safeguards
Significant Automated Decision-Making	Systems affecting consequential legal or social outcomes for individuals	ADMT, ADM systems, or equivalent requiring notice or impact assessment	Non-decision-making AI systems	California ADMT regulations: requirements concerning significant automated decisions, including notice and related assessment duties
Reasonable Care Standard	Duty-based compliance model supported by risk-management and impact-assessment obligations	Legal duties referencing NIST, ISO, or equivalent frameworks	Strict-liability or prescriptive-only regimes	Colorado AI Act: developers and deployers shall exercise reasonable care, supported by risk-management and impact-assessment duties
Agentic System Classification	Regulatory treatment of autonomous or self-directed AI systems	Explicit references to autonomous, agentic, or self-operating AI	Traditional rule-based or static models	Singapore framework: guidance on agentic systems capable of autonomous task decomposition and execution, with emphasis on adaptive oversight