



Security Control Coverage and Maturity Across Regulatory Compliance Frameworks for Generative AI Systems

Felix S. Mutau

Corporate Institute of Strategic Research. Email: samutau@gmail.com. ORCID ID: <https://orcid.org/0009-0001-7144-9089>

Received:- 18/06/2026, Revised:- 15/07/2026, Accepted:- 09/09/2026, Published:- 25/09/2026

ABSTRACT

Security governance is needed that links security technical controls to security regulatory requirements to support Generative artificial intelligence (GenAI). This study has compared the coverage and maturity of security controls as outlined in HIPAA, GDPR, the EU AI Act and financial regulations. A curated dataset of 47 requirements, 24 controls, coverage ratings, and a control–requirement mapping matrix was analyzed descriptively as a secondary analysis. One of the links was not consistent, and was not included in the identification of valid relationships, leaving 128 valid relationships. In total, 41 requirements (87.23%) were mapped with at least one control and six requirements were unmapped. Of the 96 control–framework assessments, 34.38% reported full coverage, 54.17% reported partial coverage, and 11.46% reported coverage of an unspecified nature. Controls were classified as production (58.33%), emerging (37.50%), or research (4.17%). There were 19 controls across all four regulatory groups and 16 requirements, which were supported by only one control. These results show significant common mapping and depth and maturity issues. They help you to have expert review on unmapped requirements, unspecific ratings, and control dependencies in a targeted way. These results reflect curated assessments, and are not to be used as a measure of organizational implementation, effectiveness or legal compliance.

Keywords: Generative artificial intelligence, regulatory compliance, security controls, control maturity, compliance mapping

1. Introduction

Content creation, decision-making and organizational workflows are becoming more and more automated with the help of generative artificial intelligence (GenAI). The adoption of this poses security and compliance issues around sensitive information, unreliable outputs, accountability and resilience. To tackle these challenges, it is essential to adopt coordinated technical security measures and governance practices, instead of just relying on model performance (Radanliev et al., 2025). However, the research on responsible GenAI also takes into account responsibility in data, models, users, and regulatory arrangements, as this shows that the oversight of GenAI systems goes beyond the component level (Raza et al., 2025). Security control coverage is a key topic in research on regulatory compliance, due to these concerns.

This discussion is crucial for privacy and data protection, as GenAI systems rely on data being gathered, manipulated, and utilized throughout various phases. Andrade et al. (2026) identified privacy, risk, and compliance as three interlinked concerns in the field of GenAI research through systematic mapping. Through the systematic mapping study, Andrade et al. (2026) identified privacy, risk and compliance as interconnected concerns in GenAI research. However, the question of trans-jurisdictional legal obligations to operational practices remains (Kumar, 2026). Comparing requirements can help shed light on the controls that affect common concerns and where additional focus is needed for financial services, general data protection, artificial intelligence governance, and healthcare.

These responsibilities can be structured by a risk management framework. The AI Risk Management Framework (AI RMF) profile on GenAI presents voluntary guidance for the design, development, use and evaluation of a system that considers trustworthiness. A literature review of GenAI risk taxonomies also underscores the importance of integrating ethical, cybersecurity and regulatory issues in a unified governance framework (Coutinho et al., 2026). A broad framework does not, however, come with the answers of whether or not a specific control meets a specific requirement. A review of documented control–requirement relationships can then help to bring governance discussions to life and facilitate transparent comparison.

The security aspect has to be dealt with throughout the entire AI lifecycle as well. Secure model management has been studied in a range of activities throughout the development, deployment, and maintenance lifecycle (Almagrabi & Khan, 2025). Research into autonomous systems explores the role of GenAI in enhancing security and resilience and the challenges it poses for security that must be addressed (Andreoni et al., 2024). Stewardship, data quality, data lineage, access control, and data lifecycle management are other pillars of data governance (Gangavarapu, 2025). These viewpoints lead to a proposal that compliance mapping should consider both the distribution of controls across multiple security categories, and the extent of their documented regulatory coverage.

The implementation of enterprise adds another layer of difference between the identification of controls and their maturity. Cloud-based GenAI research fosters workflow automation, secure data governance, and deployment architecture (Saravanan, 2025). A suggested enterprise governance maturity model outlines moving from experimentation to standard deployment to broader organisation optimization (Modi, 2026). The governance frameworks work also helps with systematically considering organizational controls and readiness (Emett et al., 2026). As a result, the priority of the implementation of these controls may vary – and may change from one production stage to another – and the development of these controls may be at a different level. Comparing these classifications with coverage ratings can be useful to look for descriptive patterns, not to make inferences of effectiveness or legality based on maturity.

Priorities are also different in institutional and sectoral settings. The government policy proposals have been linked to GenAI and agentic AI, regulation and oversight, and risk management (Joshi, 2025a). When discussing digital mental health devices, safety, transparency, and post-deployment support are key points that have been highlighted in proposals (Joshi, 2025b). Other institutional policies have also been analyzed with regard to cyber-insurance governance, such as acceptable use, auditing, and reporting of incidents (Hadi et al., 2026). The debate on the relationship between risk, compliance and trust is being discussed at a wider governance level (Khanna, 2026). These views highlight the importance of making cross-sectoral comparisons on the one hand, but keeping differences between regulatory contexts on the other. This study builds on the existing material, as it is accompanied by a regulatory mapping framework and is therefore described in a secondary manner. It is its part to validate the recorded links, compare the coverage profiles, analyse the common controls, and highlight requirements that are not fully mapped.

Coverage and maturity classifications are considered as assessments made by the contributor of the dataset. A

mapping does not necessarily prove implementation, effectiveness or satisfaction of a whole legal obligation, but only a documented relationship. Likewise, blank cells on the matrix do not represent violations of the organization, but rather the lack of information from the provided matrix. These distinctions help to preserve a defensible interpretation of the secondary data. The study will therefore seek to present a repeatable description of the extent of coverage and the mapping constraints that will guide future control prioritisation, expert review and empirical testing. Its objectives are:

- 1.To compare security control coverage across regulatory frameworks relevant to generative artificial intelligence systems.
- 2.To assess control maturity and examine its relationship with coverage classifications across regulatory frameworks.
- 3.To identify mapping gaps, shared controls, and requirements with limited control support across regulatory frameworks.

2. Methodology

2.1. Research Design and Dataset Description

The dataset on the regulatory compliance mapping for generative AI security (Adnan, 2025) was used to conduct a descriptive comparative analysis. The workbook consisted of three interrelated sheets: regulatory requirements, technical controls, and control-requirement mappings. Requirements and controls were seen as documentary units rather than observations from organizations. The analyzes looked at the HIPAA, GDPR, EU AI Act and financial classifications that were provided. Financial entries included Sarbanes–Oxley requirements and PCI DSS requirements. Source coverage and maturity labels remained author assessments, with no assumption of independently verified legal compliance.

2.2. Data Validation and Preparation

Data validation checked identifier uniqueness, blank cells, requirement references, and consistency of regulatory assignments between sheets. The mapping entries were broken down to individual control–requirement links and validated against the requirements register. Links assigned to an incompatible regulatory group were excluded from the primary analysis, but the original records were retained. Blank coverage assessments were not recoded as failure or nonapplicability but left as unspecified. Discrepancies in documentation were documented separately. Unresolved entries did not create a replacement requirement, and requirements not in the register were not added.

2.3. Regulatory Requirement and Control Classification

Regulatory requirements were grouped by the labels of the frameworks they provided and the frequencies of the frameworks were reported as percentages of all listed requirements. Controls were grouped into six categories: Data Minimization, Access Control, Audit and Monitoring, Data Subject Rights, Transparency, and Security. Maturity classes were grouped as Production, Emerging, or Research. The category distributions were cross-tabulated with the maturity counts to describe the balance of established and developing controls. The classifications were the judgment of the dataset authors; the analysis did not independently determine the effectiveness of implementation, the prevalence of deployment, or the acceptance of regulation.

2.4. Coverage and Mapping Analysis

Coverage assessments for each framework were summarized as Full, Partial, or Unspecified. All controls listed were used as a denominator for percentages and blank assessments were kept in the total. The binary matrix represented validated control-requirement links without assigning coverage strength to individual links. Mapping analyzes were performed to determine requirements mapping to at least one control, link totals, controls across multiple frameworks, and average links per control. The density of the matrix was computed for all possible combinations of requirement and control. The framework coverage ratings and the requirement mapping proportions were reported separately since they measured different characteristics.

2.5. Maturity, Gap, and Sensitivity Assessment

The gap analysis identified requirements that did not have validated control links and described their distribution across the frameworks. Requirements with only one mapped control were counted to indicate limited mapping redundancy, without interpreting this as insufficient security. Sensitivity analysis compared total number of links before and after validation, and tested whether the number of mapped requirements changed with exclusion. Coverage and maturity summaries were computed without reference to mapping validation. Descriptive frequencies, percentages, and network summaries were performed in Python. Inferential tests were not used as the dataset was a curated register and not a probability sample.

3. Results

3.1. Dataset Composition and Validation Findings

The workbook incorporated 47 regulatory requirements and 24 technical controls. The GDPR accounted for 18 requirements (38.30%), followed by HIPAA with 12 (25.53%), the EU AI Act with nine (19.15%) and financial regulations and standards with eight (17.02%). The composition is shown in Table 1. Identifiers were unique between the three sheets. Eleven coverage cells and eleven mapping cells were vacant. One mapping referenced a GDPR requirement in the EU AI Act column. The README also referenced Article 50, but there was no matching entry in the requirements sheet.

Table 1. Distribution of requirements across regulatory frameworks

Regulatory framework/group	Requirements, n	Percentage (%)
HIPAA	12	25.53
GDPR	18	38.30
EU AI Act	9	19.15
Financial regulations and standards	8	17.02
Total	47	100.00

3.2. Security Control Categories and Maturity

Each of the six control categories had four controls, or 16.67 percent of the control register. Of the 24 controls, 14 (58.33%) were classified as Production, nine (37.50%) as Emerging, and one (4.17%) as Research. Table 2 presents the maturity distribution by category. Access Control had four Production controls and the other four categories had two Production and two Emerging controls each. Data Subject Rights included two Production, one Emerging and one Research control. The only Research control was Machine Unlearning which indicated the source classification and was not independently assessed for readiness (Figure 1).

Table 2. Control maturity distribution across six security categories

Control category	Production, n	Emerging, n	Research, n	Total, n
Data Minimization	2	2	1	5
Access Control	4	1	1	6

Audit and Monitoring	2	2	1	5
Data Subject Rights	2	1	1	4
Transparency	2	2	1	5
Security	2	2	1	5
Total	14	10	6	30




Figure 1. Average requirement mapping breadth across security control categories

3.3. Comparative Coverage Across Regulatory Frameworks

Out of 96 control-framework assessments, 33 (34.38%) were labeled Full, 52 (54.17%) Partial, and 11 (11.46%) were unspecified. Table 3 shows distributions specific to framework. The largest Full proportion was with GDPR with ten controls (41.67%), followed by the EU AI Act with nine (37.50%). There were seven Full assessments of both HIPAA and financial entries (29.17%). Mostly all frameworks were partial in the assessments. HIPAA had the highest number of unspecified assessments (5, 20.83%). Of the 33 Full assessments, 26 (78.79%) were production controls, which did not show better practical effectiveness (Figure 2).

Table 3. Numerical comparison of framework-level control coverage assessments

Regulatory framework/group	Full, n (%)	Partial, n (%)	Unspecified, n (%)
HIPAA	7 (29.17)	12 (50.00)	5 (20.83)
GDPR	10 (41.67)	12 (50.00)	2 (8.33)
EU AI Act	9 (37.50)	14 (58.33)	1 (4.17)
Financial regulations and standards	7 (29.17)	14 (58.33)	3 (12.50)
All assessments	33 (34.38)	52 (54.17)	11 (11.46)




Figure 2. Coverage assessment profiles across security control maturity levels

3.4. Control–Requirement Relationships and Shared Controls

Validation retained 128 control–requirement links. These linked 41 of 47 requirements (87.23%) to at least one control and ended in a matrix density of 11.35%. Table 4 summarizes the mapping coverage for each framework. Nineteen (79.17%) controls had links across all four frameworks. Federated Learning had the highest number of links, addressing eight requirements, followed by AI Disclosure and Deepfake Labeling with one requirement each. HIPAA: 33 links GDPR: 39 links EU AI Act: 28 links Financial group: 28 links Average validated links per control was 5.33 (Figure 3).

Table 4. Validated requirement mapping coverage across regulatory frameworks

Regulatory framework/group	Requirements, n	Mapped requirements, n (%)	Unmapped requirements, n	Validated links, n
HIPAA	12	11 (91.67)	1	33
GDPR	18	14 (77.78)	4	39
EU AI Act	9	8 (88.89)	1	28
Financial regulations and standards	8	8 (100.00)	0	28
Total	47	41 (87.23)	6	128




Figure 3. Validated requirement links across individual security controls

3.5. Unmapped Requirements and Sensitivity Findings

Six requirements (12.77%) did not have validated control links: HIPAA Contingency Planning; GDPR Purpose Limitation, Right to Rectification, Right to Restriction and DPIA Requirement; and EU AI Act Conformity Assessment. All financial requirements were mapped as shown in Table 4. Sixteen requirements (34.04%) had a single control mapped. Removing the inconsistent Confidence Scores–R13 entry reduced links from 129 to 128 but did not change mapped-requirement coverage, which remained at 87.23%. Coverage and maturity counts were not impacted. What we found were gaps in the mapping that was provided, not instances of non-compliance by the organization (Figure 4).




Figure 4. Distribution of mapped requirements by control support

4. Discussion

A wide geographic area can be covered with shallow coverage. While the 41 of 47 requirements had at least one validated control link, 34.38% of the 96 control–framework assessments were shown as full coverage and 54.17% were partial. Such behavior is in line with the literature, which highlights challenges in bridging between the identified AI-specific vulnerabilities, technological capabilities, and security conformity assessment with regulatory expectations (Hamon et al., 2024). Documenting a relationship is a first step to compliance analysis, while showing adequate protection involves looking at implementation conditions, residual risks, and the specific obligation being addressed.

There is significant intersection among the various framework-specific control sets that support literature that calls for integrated trust, risk, and security management across the entire AI lifecycle (Ray, 2026). All four frameworks were connected to 19 controls and Jaccard similarities were between 82.61% and 95.45% in pairs. The findings suggest opportunities to structure shared safeguards in a shared governance structure. But, there are no legal obligations or implementation needs which are equal if control sets overlap. The matrix is now more limited than the TRiSM overview in that it only identifies some documented relationships, and does not identify whether shared controls provide sufficient protection in other operational contexts.

Control maturity also was unevenly distributed with 58.33% production maturity, 37.50% emerging maturity, and 4.17% research maturity. The proportion of assessments that were full coverage was 46.43%, while the proportion of assessments in emerging-control was 19.44%. This descriptive difference aligns with governance literature that relates organizational capability assessment to lifecycle practices, auditing and rights respecting implementation (Almeida, 2024). However, the maturity level labels for the controls are not the same as the one that is applicable to the governance maturity at the organisation level. The pattern observed, therefore, does not justify a causal statement that production maturity is what leads to greater compliance, but rather differentiated review priorities. No outcomes of implementation and organizational readiness were available.

Data minimisation controls had more mapped relationships than transparency controls, with an average of 6.75 mapped controls per control compared to 2.75. There was one validated link for each of AI disclosure and deepfake labelling, and six for data provenance. This comparison is relevant to literature that suggests a technical detection approach should be leveraged alongside policy to combat misinformation and promote trustworthy governance of GenAI (Abbas et al., 2026). The lower mapping breadth does not mean that the obligations are less significant; specialized transparency controls can accommodate narrowly defined obligations. The literature has a more general perspective on society, while the dataset only reflects certain regulatory relationships.

In the access-control schema, the four measures were all classified as production; this is an area of data in which there has been higher consistency of maturity. This pattern is conceptually similar to the layered enterprise security pattern that focuses on identity governance, least privilege, telemetry, guardrails, and enforcement at agentic decision boundaries (Meesala, 2023). But existing access controls are just part of the defense. The present analysis does not include any measurement of detection latency, containment time, or audit readiness as would be the case in an operational assessment as described in that framework. This means that the comparison can be considered alongside access control and monitoring and model protection, although there is no claim to operational effectiveness.

Compared to cloud literature, which combines disaster recovery and business continuity with secure infrastructure to support AI capabilities (Pandey & Patel, 2026), the unmapped HIPAA contingency-planning requirement highlights the crucial need for safeguarding sensitive patient data. The unmapped HIPAA contingency-planning requirement stands out as a key element in protecting sensitive patient data, while cloud literature focuses on integrating disaster recovery and business continuity with secure infrastructure to support AI capabilities. This apparent completeness did not permit getting rid of partial or unspecified coverage assessments, but all financial requirements were mapped with controls. The findings show the importance of the separate evaluation of the completeness of the mapping and the sufficiency of the implementation of the framework. The matrix may be seen as a representative selection of compliance support, as compared to the various aspects of infrastructure, governance and continuity highlighted in the cloud literature. The contingency-planning gap is thus not a reflection of the absence of contingency planning arrangements in healthcare organizations, but rather represents a documentation priority in this dataset.

Another point of comparison in healthcare is the context of deployment. While HIPAA had 11 of its 12 mapped requirements, it also had 12 partial ratings and five unspecified entries in its coverage assessments. Policy literature in healthcare suggests a tiered governance, which ought to be responsive to the variety of characteristics that can be

found in open source, proprietary and agentic models (Joshi, 2025c). These deployment arrangements are not distinguished here in the present dataset as they would be in that proposed framework and so cannot be compared directly with the present dataset. However, the breadth of mapping and partial ratings are aligned with the requirement for assessment to be context sensitive. Beyond mapped controls, Privacy safeguards, accountability arrangements and validation evidence need to be evaluated.

Lastly, 16 of the requirements had only one mapped control, indicating a lack of documented redundancy. The literature on GenAI safety and security calls for a multi-layered sociotechnical assessment that includes human interactions, governance processes, and system conditions as well as the capability of the models (Janjeva et al., 2023). A single control mapping provides minimal evidence of resilience in the event of a safeguard failure or failure to apply the safeguard, compared with that approach. Several links cannot provide independent security: controls can depend on one another. The comparison is then more likely to be against an expert view of the control interactions and operational evidence and the present results are only descriptive and informative about the provided compliance matrix.

5. Conclusion

Security control coverage and maturity is defined for a set of 47 security requirements and 24 security controls across four regulatory frameworks. Following validation, 41 requirements (87.23%) were mapped with one or more controls and 128 valid links were created. Full coverage, however, was only observed in 34.38% of control–framework assessments, 54.17% of which were partial and 11.46% of which were not specified. Production controls made up 58.33% of the control inventory, with 37.50% of controls emerging and 4.17% at a research stage. The percentages of full coverage ratings were greater for production controls, but this is merely a descriptive pattern and does not imply a causal relationship between maturity and compliance. A large degree of overlap between frameworks suggests opportunities to align common safeguards in a unified governance model. However, six requirements that were not mapped out and 16 requirements with only one control indicate there are areas requiring further investigation. The results align with the idea that mapping completeness, coverage depth and control maturity are three distinct aspects that should be considered separately when prioritizing generative AI security governance. Specification of ratings; review of requirements not explicitly stated; verification of mapping accuracy by independent expert evaluation; assessment of functionality of linked controls in operational deployments should be resolved in subsequent assessment. The results of this database are evidence of documented coverage patterns and gaps in mapping, but not evidence of implemented safeguards or regulatory compliance or security effectiveness.

References

1. Abbas, F., Chesterman, S., & Taeiagh, A. (2026). Building trust in the generative AI era: a systematic review of global regulatory frameworks to combat the risks of mis-, dis-, and mal-information. *AI & SOCIETY*, 41(4), 3913-3950.
2. Adnan, M. (2025). Supplementary data: Regulatory compliance mapping for GenAI security—Requirements, controls, and mapping matrix (Version 1) [Data set]. Mendeley Data. <https://doi.org/10.17632/mm6tfm663v.1>
3. Almagrabi, A. O., & Khan, R. A. (2024). Optimizing secure AI lifecycle model management with innovative generative AI strategies. *IEEE Access*, 13, 12889-12920.
4. Almeida, I. (2024). Responsible AI in the Age of Generative Models: Governance, Ethics and Risk Management. *Now Next Later AI*.
5. Andrade, R. B. D. S., Rocha Filho, G. P., dos Santos Marques, G., & Canedo, E. D. (2026). Privacy, Data Protection, Risk, and Compliance in the Age of Generative AI Systems: A Systematic Mapping Study. *Simpósio Brasileiro de Sistemas de Informação (SBSI)*, 1181-1200.
6. Andreoni, M., Lunardi, W. T., Lawton, G., & Thakkar, S. (2024). Enhancing autonomous system security and resilience with generative AI: A comprehensive survey. *IEEE Access*, 12, 109470-109493.
7. Coutinho, M. A., Ashofteh, A., & Al Helaly, Y. (2025, June). Risk taxonomies and governance frameworks for generative AI: A review of ethical, cybersecurity, and regulatory challenges. In *Iberian Conference on Information Systems and Technologies* (pp. 3-15). Cham: Springer Nature Switzerland.
8. Emmett, S., Eulerich, M., Pikoos, J., & Wood, D. A. (2026). The development of a generative artificial intelligence

- (AI) governance framework. *Accounting Horizons*, 1-16.
9. Gangavarapu, R. (2025). Data Governance as an Enabler for Generative AI Maturity. In *Mastering AI Governance: A Guide to Building Trustworthy and Transparent AI Systems* (pp. 35-43). Cham: Springer Nature Switzerland.
 10. Hadi, W., Hamarsheh, A., Al-Qerem, A., & Aldweesh, A. (2026). Institutional Policies for Generative AI Implications for Cyber-Insurance Governance. In *Cross-Sector Cyber Insurance for the Intelligent Society* (pp. 83-112). IGI Global Scientific Publishing.
 11. Hamon, R., Junklewitz, H., Garrido, J. S., & Sanchez, I. (2024). Three challenges to secure AI systems in the context of AI regulations. *Ieee Access*, 12, 61022-61035.
 12. Janjeva, A., Harris, A., Mercer, S., Kasprzyk, A. M., & Gausen, A. (2023). The rapid rise of generative AI: Assessing risks to safety and security.
 13. Joshi, S. (2025a). Framework for government policy on agentic and generative AI: Governance, regulation, and risk management.
 14. Joshi, S. (2025b). Regulatory frameworks for generative AI enabled digital mental health devices: safety, transparency, and post-market oversight. *Transparency, and Post-Market Oversight* (December 08, 2025).
 15. Joshi, S. (2025c). Framework for government policy on agentic and generative AI in healthcare: Governance, regulation, and risk management of open-source and proprietary models. *Regulation, and Risk Management of Open-Source and Proprietary Models* (July 02, 2025).
 16. Khanna, D. (2026). *Governing Intelligence: Strategies for Managing Risk, Compliance, and Trust in the Age of Generative AI*. Geh press.
 17. Kumar, A. (2026). Legal and regulatory frameworks governing generative AI for enterprises. In *GenAI and LLMs for beyond 5G Networks* (pp. 39-81). Cham: Springer Nature Switzerland.
 18. Meesala, L. K. (2023). A layered security framework for enterprise operations in the Generative AI and Agentic AI era in regulated cloud environments. *International Journal of Future Innovative Science and Technology (IJFIST)*, 6(6), 11760.
 19. Modi, R. (2026). An Enterprise AI Governance Maturity Model for Cloud-deployed Generative AI Systems. Available at SSRN 6624698.
 20. Pandey, B., & Patel, A. (Eds.). (2026). *Revolutionizing the Cloud: Generative AI, Security, and Sustainability*. Springer Nature.
 21. Radanliev, P., Santos, O., & Ani, U. D. (2025). Generative AI cybersecurity and resilience. *Frontiers in Artificial Intelligence*, 8, 1568360.
 22. Ray, P. P. (2026). A review of TRiSM frameworks in artificial intelligence systems: Fundamentals, taxonomy, use cases, key challenges and future directions. *Expert Systems*, 43(3), e70213.
 23. Raza, S., Qureshi, R., Zahid, A., Muneer, A., Zafar, A., Kamawal, S., ... & Shoman, M. (2025). Who is responsible? the data, models, users or regulations? a comprehensive survey on responsible generative ai for a sustainable future. *arXiv preprint arXiv:2502.08650*.
 24. Saravanan, K. (2025). Cloud-Native Generative AI Frameworks for Enterprise Workflow Automation and Secure Data Governance. *International Journal of Computer Technology and Electronics Communication*, 8(6), 11952-11961.