

The Autonomous Compliance-to-Code (AC2C) Framework: A Generative AI-Powered Paradigm for Mapping IoT Data Flows to GDPR, HIPAA, and NIS2

Adetunji Oludele Adebayo¹, Omowunmi Folashayo Makinde², Olatunde Ayomide Olasehan³, Nathaniel Adeniyi Akande⁴, & Udoka Cynthia Duruemeruo⁵

¹ Information Security Manager /Independent researcher, University of Bradford, UK

²IT Support Engineer I/Independent Researcher, University of the Cumberland, US

³IT Engineer/Independent Researcher, Swansea University, UK

⁴Cybersecurity Analyst/Independent Researcher, University of Bradford, UK

⁵DevOps Engineer/Independent Researcher, University of Wolverhampton, UK

DOI - <http://doi.org/10.37502/IJSMR.2025.81208>

Abstract

The proliferation of Internet of Things (IoT) devices has transformed industries, including healthcare, by enabling real-time monitoring, optimization, and predictive maintenance. However, this advancement has introduced complex regulatory requirements, such as the General Data Protection Regulation (GDPR), the Health Insurance Portability and Accountability Act (HIPAA), and the Network and Information Systems Directive 2 (NIS2). Traditional governance, risk, and compliance (GRC) models have proven insufficient in addressing the dynamic nature of IoT deployments, leading to a compliance gap. This paper introduces the Autonomous Compliance-to-Code (AC2C) framework, a generative AI-powered paradigm for mapping IoT data flows to GDPR, HIPAA, and NIS2. The AC2C framework uses a multi-agent architecture to interpret regulatory requirements, map them to a dynamic technological environment, synthesize executable compliance logic, and perform continuous assurance. The framework's efficacy is demonstrated through a case study involving a global smart healthcare provider subject to multiple jurisdictions. The framework's agents, including the Regulatory Deconstruction Agent (RDA), IoT Data Flow Intelligence Agent (DFIA), Compliance Logic Synthesis Agent (CLSA), and Continuous Assurance & Reporting Agent (CARA), work together to bridge the "regulation-to-code" bottleneck and enable proactive compliance in IoT ecosystems. The AC2C framework represents an advancement in RegTech, offering organizations an automated solution to navigate IoT compliance.

Keywords: Compliance-to-Code, GDPR, HIPAA.

1. Introduction

The evolution of digital critical infrastructure has dramatically changed the compliance climate in all industries. The Internet of Things (IoT) has become a game-changer, with more than 15 billion enabled devices installed worldwide as of 2023. This phenomenon impacts key industries, ranging from healthcare to the energy sector (Zaman, 2023). This proliferation has spawned groundbreaking applications that span real-time patient monitoring systems,

intelligent grid optimization, and predictive maintenance operations which serve to raise levels of operational efficiency and service quality (Junaid et al., 2022).

But this same technological development has also given rise to a complex and cumbersome array of onerous regulatory requirements that stretch around the world and well beyond existing frameworks. Healthcare IoT deployments require compliance with the stringent regulations enforced by the Health Insurance Portability and Accountability Act (HIPAA) in the United States, with European operations subject to the stringent general data protection regime of the General Data Protection Regulation (GDPR) (Said et al., 2024). The new arrival has now entered the scene with the Network and Information Systems Directive 2 (NIS2) which imposes new cybersecurity obligations on essential and important entities located in the European Union (Chiara, 2022).

Conventional governance, risk and compliance (GRC) models designed for static, centralized IT have been insufficient to address the dynamic, decentralized nature of IoT deployments. These techniques are based primarily on regular audits, manual documentation checks, and static policy models that are not capable of keeping pace with the rapid change of technology or the uninterrupted data traffic inherent to IoT deployments. The outcome is an increasingly widening gap in compliance that is putting at risk institutions with substantial fines, operational interruptions and reputational harm (Halgamuge & Niyato, 2025).

At the center of this compliance crisis is what we call the "regulation-to-code" bottleneck, which is a fundamentally knowledge translation gap and the most significant impediment to successful regulatory technology implementation. The bottleneck refers to the manual, labor-intensive, and error-prone process of interpreting vague legal language, and translating that into accurate and executable compliance logic. Compliance-as-Code (CaC) models, while offering a considerable advancement in the scale of GRC, are still hampered by needing human intelligence to close the meaningful gap between regulatory intent and technical implementation (Sardana et al., 2024). Legal professionals have a deep understanding of compliance obligations, but lack the expertise to articulate them in a way that can be converted to a machine-readable policy. Technical teams understand system architecture and coding frameworks, but struggle to interpret the nuanced language of regulatory documents and their implications (Li et al., 2023).

This translation complication is complicated by uncertainty in regulatory language that is often intentionally vague, allowing for various organizational contexts and technology applications. Phrases such as "appropriate technical measures," "reasonable safeguards," and "proportionate response," require contextual interpretation and can differ greatly depending on the size of the organization, the sector it operates in, the geographic location, and the infrastructure within which the technology operates. Consequentially, the issues presented in this bottleneck go beyond issues of compliance inefficiency. Organizations are increasingly facing expectations from regulators who desire compliance monitoring in real-time and responses to incidents in a timely manner (Schmitz-Berndt, 2023).

In the case of the European Union, there are 72-hour notice requirements for breach notifications stated in the GDPR (General Data Protection Regulation), and NIS2 requires 24-hour early warning notifications for significant incidents. Such demanding timelines cannot be accomplished through outdated and inefficient manual processes. The end result is an

uncontrolled mismatch between the regulatory environments and the capacities of organizations (Chiara, 2022).

In addition, the agile nature of IoT environments compounds these issues. IoT solutions are very dynamic and involve rapid introduction of new devices, firmware updates, changes in data streaming, and integration of new services. Every single change has the potential for impacting compliance posture, and therefore must be reassessed on an ongoing basis for regulation compliance and control effectiveness. Manual processes will not scale to meet these new challenges, causing long-term blind spots and gaps in compliance coverage (Hornos & Quinde, 2024). The financial burden is quite considerable. Organizations indicate that they spend 15-20% of their IT budgets on compliance activities, with most of that money directed on manual processes that have seldom offered any continued value. Adding the opportunity cost of shifting technical resources around from innovation to compliance administration are just added costs to gravity of the original costs (Ghafari et al., 2024).

In this paper, we are introducing a new generative Artificial Intelligence (AI)-powered framework for compliance called Autonomous Compliance-to-Code (AC2C) Framework, which aims to remove the regulation-to-code bottleneck, and help organizations intelligently automate all aspects of the compliance lifecycle, from regulatory interpretation, to continuous assurance.

2. Evolution of Regulatory Technology (RegTech)

Regulatory Technology (RegTech) is a significant change in the manual and spreadsheets processes to an automated tool to governance, risk and compliance (GRC). Compliance was handled manually or by crudely developed digital tools in the early days. However, as global regulations became more complex and stringent, this approach proved insufficient. The rise of Compliance-as-Code (CaC) is a direct result of this need for more efficient and scalable compliance solutions. Inspired by Infrastructure-as-Code (IaC), CaC allows organizations to codify regulatory policies into machine-readable formats, enabling automated compliance validation and real-time monitoring (McKinsey & Company, 2020). This is essential in any industry where compliance has to be strictly observed on a continuous basis, like in finance and healthcare where a breach of regulations would result in substantial fines. With the help of compliance processes turned into automated systems, organizations can make sure that compliance is a part of the operations workflow with the possibility of real-time auditing and risk management. It is worth noting that this automation has contributed to the scaling of GRC systems, which have minimized dependence on human resources and minimized the chances of mistakes. Nevertheless, even with these developments, regulatory frameworks continue to develop and systems need to be dynamic. As a result, RegTech has become a central pillar in the digital transformation of compliance functions, aiming to address the shortcomings of legacy systems and provide businesses with tools that can respond to dynamic regulatory environments more effectively.

2.1 Compliance Challenges in the RegTech Landscape

Despite the significant advancements in Regulatory Technology, numerous challenges persist in the field of Compliance-as-Code (CaC). One of the most significant barriers is the so-called “regulation-to-code bottleneck,” which refers to the inherent difficulty in translating legal language into machine-readable formats. Regulatory texts, such as the General Data Protection

Regulation (GDPR) and the Health Insurance Portability and Accountability Act (HIPAA), are often ambiguous and complex, requiring specialized legal knowledge to interpret accurately (European Commission, 2020). This activity has traditionally been the role of legal teams, and RegTech solutions fail to automate it effectively, given the complexity of legal language, which is not always readable by the existing AI technology. In this regard, even though RegTech tools such as CaC are automated, they continue to require a high degree of human involvement in terms of translation and implementation. Moreover, regulatory frameworks do not remain consistent; they change regularly, and compliance tools have to follow suit. Conventional systems based on hard-codes can soon become obsolete and this poses a trap of inflexibility, which makes them expensive as well as challenging to support. This is particularly evident in rapidly changing sectors such as healthcare and finance, where continuous compliance is crucial. Therefore, RegTech faces the dual challenge of overcoming the bottleneck of legal interpretation and ensuring flexibility to accommodate ongoing regulatory changes.

2.2 Philosophical Mismatch in Traditional Compliance Tools

A significant issue in traditional compliance tools lies in the philosophical mismatch between the “auditing” mindset of many existing systems and the “threat-hunting” approach required by modern regulations like NIS2 (European Commission, 2023). Traditional compliance systems focus on auditing and ensuring that organizations meet regulatory requirements at a given point in time, often through manual or periodic checks. However, as cybersecurity threats evolve, regulations like NIS2 demand a more proactive approach. This regulatory shift requires tools that can constantly monitor, detect, and mitigate risks in real time, rather than simply auditing compliance after the fact. For instance, under NIS2, organizations are expected to adopt a more dynamic approach to security, requiring continuous monitoring and response to potential threats. This is a challenge to conventional compliance systems which are very inflexible and not meant to accommodate such proactive requirements. The contrast between the two thinking approaches is that one is backward-looking and the other forward-looking that points out a basic deficiency in the existing regulatory frameworks. Therefore, the need for automated compliance tools that can continuously adapt to changing security landscapes becomes evident. In this context, technologies like Generative AI show promise in bridging this gap by automating both the compliance process and the monitoring required to meet these dynamic regulatory requirements (NIST, 2025).

2.3 Role of Generative AI in Compliance

Generative AI has the potential to fundamentally transform how compliance processes are managed by automating the interpretation and generation of compliance logic. As previously mentioned, the regulation-to-code bottleneck remains a major challenge in RegTech. Generative AI, especially Large Language Models (LLMs), offers a promising solution by automating the extraction of legal intent from regulatory texts and translating that into actionable compliance logic (Brown et al., 2020). This ability to generate compliance rules from unstructured legal texts is particularly beneficial in contexts where regulations are complex and frequently changing. Moreover, LLMs can process vast amounts of data quickly, enabling continuous compliance validation without the need for manual intervention. These AI models can also be fine-tuned to specific regulatory frameworks, such as GDPR, HIPAA, and NIS2, making them adaptable across industries. As an example, AI can understand a regulatory language, and base its interpretations on authoritative legal sources to ensure soundness using

Retrieval-Augmented Generation (RAG) models. The fact that these models are real-time also contributes to constant monitoring and the ability to adjust to the requirements of compliance, which is especially useful in finance and healthcare sectors, where the regulations change fast. However, while Generative AI holds great promise, its application in compliance must be handled carefully to avoid misinterpretation of legal language, as AI models can still struggle with the subtleties of legal jargon (Gartner, 2024).

2.4 Gaps and Future Directions

Despite the transformative potential of Generative AI in the compliance space, significant gaps remain that must be addressed. One of the primary challenges is the translation accuracy of legal texts. While Generative AI models can assist in interpreting regulations, they are not immune to errors, particularly when faced with ambiguous or contradictory legal language (Sadiku, Ajayi and Ajayi, 2025). Moreover, such models demand considerable volumes of good-quality legal information to be trained, which may be a constraint in some jurisdictions where such information is not readily available or may be hidden. A second gap is that there is a necessity to constantly adapt due to the changing legal frameworks. There should be a seamless integration of regulatory changes into AI systems in order to maintain compliance. The existing systems are usually unable to deal with this constant evolution and thus a delay in compliance is experienced when the regulations are altered. Future research should focus on improving the interpretative capabilities of AI models and ensuring that these systems are capable of handling new types of regulations that may emerge (Mucci, 2024). Additionally, ongoing work is required to reduce the dependency on human intervention for the fine-tuning of AI models, thereby enhancing the scalability and efficiency of automated compliance systems.

3. Methodology

This study introduces and evaluates the Autonomous Compliance-to-Code (AC2C) framework, a novel, closed-loop, multi-agent Generative AI system designed to automate the compliance lifecycle for complex Internet of Things (IoT) ecosystems (Li et al.). The research methodology follows a constructive approach, wherein a new technological artefact, the AC2C framework, is designed, implemented in a detailed case study, and evaluated based on its capacity to solve a real-world problem (Li et. al).

3.1 Research Design

The research is designed as a qualitative, descriptive case study. The central artefact is the AC2C framework, which employs a multi-agent architecture to continuously interpret regulatory requirements, map them to a dynamic technological environment, synthesise executable compliance logic, and perform continuous assurance (Ettaloui et al.). The efficacy of this design is demonstrated through its application to a realistic, multi-jurisdictional scenario involving a global smart healthcare provider subject to the General Data Protection Regulation (GDPR), the Health Insurance Portability and Accountability Act (HIPAA), and the NIS2 Directive. This approach enables an in-depth examination of the framework's capabilities in a complex, real-world context.

3.1.1 Framework Architecture and Procedures

The AC2C architecture is composed of four distinct, interconnected agents that operate in a continuous cycle of interpretation, mapping, synthesis, and verification (Uhrmacher et al., 2024, Wimmer et al., 2009; Rullmann et al., 2007; Nabipour et al., 2025; Lee et al., 2008).

3.1.1.1 Regulatory Deconstruction Agent (RDA)

The RDA is designed to transform unstructured legal and regulatory texts into a structured, machine-understandable format. The agent employs a Retrieval-Augmented Generation (RAG) model. The base Large Language Model (LLM) is fine-tuned on a specialised corpus of legal, cybersecurity, and technical documents to enhance its domain-specific comprehension (Barron et al., 2025). When processing a compliance query, the RAG model first retrieves relevant clauses from the source texts to ground its reasoning and mitigate factual inaccuracies (Shen et al.). It then interprets these passages to construct a multi-relational Compliance Knowledge Graph, which formally represents the regulatory landscape, its core principles, obligations, and logical relationships.

3.1.1.2 IoT Data Flow Intelligence Agent (DFIA)

The DFIA provides a real-time, dynamic map of the organisation's technological environment. Adopting a continuous discovery methodology, the agent integrates with diverse data sources, including network traffic analysis tools, cloud provider APIs, device manifests, and Software Bills of Materials (SBOMs). This process moves beyond static, interview-based data mapping exercises (Wu et al.) to produce and continuously maintain a live model of the IoT data flow and service dependency landscape, capturing asset inventories, data classifications, storage locations, and third-party dependencies.

3.1.1.3. Compliance Logic Synthesis Agent (CLSA)

The CLSA serves as the core of the framework, designed to autonomously bridge the "regulation-to-code" bottleneck. It fuses the structured Compliance Knowledge Graph from the RDA with the real-time Data Flow Map from the DFIA. Powered by a state-of-the-art code-generation LLM, the CLSA reasons over these inputs to synthesise context-specific, executable compliance tests and policies in languages such as Python or Open Policy Agent (OPA)'s Rego. This logic is generated on demand, tailored precisely to the observed state of the environment and the applicable regulatory requirements.

3.1.1.4. Continuous Assurance & Reporting Agent (CARA)

The CARA operationalises the framework by functioning as a continuous auditing and assurance engine (Moon & Krahel). It executes the suite of tests generated by the CLSA against the live environment at high frequency. Upon detecting a test failure, the CARA generates a real-time, context-enriched alert that traces the violation back to the specific regulatory clause from the RDA's knowledge graph (Alles et al.). This process emphasises explainability, a key principle in AI governance (National Institute of Standards and Technology [NIST], n.d.). The agent produces time-stamped, auditable reports that serve as immutable evidence of the organisation's compliance posture.

3.1.2 Case Study Application

To validate the framework, a detailed case study of "HealthSphere," a global smart healthcare provider, was conducted. The scenario was designed to test the AC2C framework's ability to

navigate the overlapping and conflicting requirements of GDPR, HIPAA, and NIS2. The framework was applied to HealthSphere's IoT ecosystem, which included wearable devices, a mobile application, and a cloud analytics platform operating in both the EU and the U.S. The actions of each agent were traced through specific compliance scenarios, including cross-border data transfers, supply chain security validation, and a simulated multi-regulatory incident response, to demonstrate the framework's practical utility and effectiveness.

4 A.C2C Framework: Agent-by-Agent

4.1 Regulatory Deconstruction Agent (RDA)

4.1.1 Functionality

The Regulatory Deconstruction Agent (RDA) plays a critical role in transforming the highly complex and often ambiguous world of legal documentation into structured, AC2C Framework: Agent-by-Agent machine-understandable knowledge that can be applied directly to compliance automation in Internet of Things (IoT) systems. In practice, this means that the RDA ingests large volumes of legal texts that are unstructured, for instance an overarching frameworks such as the General Data Protection Regulation (GDPR) to more domain-specific rules like the Health Insurance Portability and Accountability Act (HIPAA), Payment Card Industry Data Security Standard (PCI DSS), and Family Educational Rights and Privacy Act (FERPA), among others and restructures them into acceptable computational systems format. This process is indispensable in the IoT landscape, where regulatory obligations are not only vast in scope but also nuanced in interpretation. Therefore, while the RDA provides a foundation to alleviate the burden of operationalizing compliance by enabling automated systems to understand and apply regulatory requirements in real time, researchers have highlighted the increase in this burden as regulators now impose fines and penalties for non-compliant entities (Kulkarni et al., 2021).

4.1.2 Methodology

The research method employs a combination of a Retrieval-Augmented Generation (RAG) architecture model, hybrid multi-agent KG pipelines, Legal-norms adaptation, and mitigating hallucination with explainability to enhance accuracy. RAG architecture leverages structured KGs and is combined with LLM to substantially improve both the accuracy and contextual relevance compared to unstructured vector retrieval alone. With the RAG model, the RDA can interpret legal texts and cite the relevant sources to mitigate hallucinations and enhance accuracy (Singh, 2024). This is also very effective when IoT data flows intersect with multiple regulatory domains simultaneously, which requires balancing when legal obligations overlap (Echenim & Joshi, 2023). Consequently, the RDA bridges the gap between the regulations and concrete system design.

4.1.3 Output

The multi-relational Compliance Knowledge Graph (CKG) is generated as the result of the RDA process, representing the regulatory landscape that helps to facilitate compliance, enhance data privacy, and security. The graph serves as a foundational tool to map the IoT data flows to the relevant regulations (Xu et al., 2024). The graph also enables the mapping from legal text to structured form, ensures citations are traced to the original regulatory clauses, and performs queryable constructs for downstream agents like CARA or DFIA. For example, it can be used

to automatically verify whether data collected by a smart health monitoring device complies with HIPAA provisions or whether cross-border data transfers are consistent with GDPR restrictions. In this way, the RDA sets the stage for proactive compliance, rather than reactive enforcement, within dynamic IoT ecosystems.

4.2 IoT Data Flow Intelligence Agent (DFIA)

The IoT DFIA is a component that introduces dynamic and heterogeneous ecosystems in which data flows across interconnected IoT devices by providing real-time insights into data flows and their related service dependencies. The central purpose is to generate real-time data transmission, processing, and storage across distributed services, thereby enhancing situational awareness and resulting in a dynamic graph that visualizes these interactions.

4.2.1 Continuous Discovery Methodology

To ensure that system monitoring is not restricted to static snapshots, the DFIA is designed and operates on the principle of continuous discovery. This is achieved through the integration of multiple data sources from the network traffic analysis, cloud APIs, and the Software Bills of Materials (SBOMs) to enhance both operational behaviour and latent risks in real time (Kilaru et al., 2024).

The Network traffic analysis provides protocol-level insights into device communications and potential anomalies, while cloud APIs expose interactions between IoT devices and cloud-hosted services. To complement the two above, the Software Bills of Materials (SBOMs) contribute valuable metadata regarding component provenance and software dependencies, enabling security teams to detect vulnerabilities at the supply-chain level.

4.2.2 Output and Visualization

The output of the DFIA is a graph that visualises IoT data flows and various service dependencies. The DFIA provides a foundation for aligning system behaviour with regulatory obligations, and achieves this by not only facilitating threat detection and response, but also ensuring policy enforcement. The graph-based visualisation aligns with some of the IoT security frameworks, such as IoTGuard and AI4SAFE-IoT, which emphasize proactive monitoring and safeguarding connected environments (Shahin et al., 2025). While it is important to acknowledge that other frameworks focus more heavily on anomaly detection or resource efficiency (Chaganti, 2025), the DFIA highlights how information traverses between devices and services.

4.2.3 Compliance Logic Synthesis Agent (CLSA)

The CLSA leverages the integration of multiple frameworks and regulations, such as GDPR, NIS2, and HIPAA, and utilizes IoT frameworks to ensure comprehensive data handling (Echenim & Joshi, 2023). To ensure that dynamic IoT data flows remain aligned with evolving governance, risk, and assurance requirements, the CLSA leverages semantic models and policy-as-code frameworks to transform regulations into machine-interpretable rules. With the CLSA, auditable compliance, automated enforcement, and continuous monitoring are enabled.

4.3 Continuous Assurance & Reporting Agent (CARA)

4.3.1 Functionality

The Continuous Assurance & Reporting Agent (CARA) acts as a real-time compliance auditor. It consumes the logic synthesized by RDA (via KG) plus data flows from DFIA, monitoring for violations, and generating immutable, time-stamped evidence. Its outputs include Real-time compliance dashboards, enriched alerts when testing thresholds fail, and auditable reports with clear trails for review and accountability. With CARA functioning as a real-time compliance auditing engine, it issues context-rich alerts whenever a compliance test fails, and this continuous assurance capability is especially vital in IoT ecosystems, where the sheer volume, velocity, and variability of data flows make traditional, periodic auditing approaches insufficient.

4.3.2 Methodology

The time-stamped audit trails generated in CARA are immutable, making every compliance test carried out and its outcomes recorded in a manner that cannot be altered retroactively. Consequently, the integrity of compliance evidence from CARA is traceable and reliable, which makes it perfectly aligned with best practices for both regulatory oversight and cybersecurity, where transparency and accountability are essential. Furthermore, the continuous auditing that CARA adopts addresses the dynamic nature of IoT systems, where data flows, device interactions, and contextual factors change in real time (Pasquier et al., 2018). Therefore, CARA provides organizations with the assurance that their IoT infrastructures remain compliant, through the constant vigilance it provides even when external regulations evolve and internal system configurations are varied.

4.3.3 Output

CARA generates a range of dashboards, alerts, and reports as outputs to support both operational decision-making and regulatory reporting. The Real-time compliance dashboards are interactive visuals that highlight key compliance metrics across different regulations, such as GDPR, HIPAA, and NIS2, among others, and provide timely breach reports. While the generated outputs via alerts are context-rich notifications that state the violated KG rule, reason for the violation, and supporting evidence such as time-stamped event logs and regulatory clause citation, the reports are time-stamped exportable PDFs or JSON logs with cryptographic hash chains guaranteeing integrity and acting as concise audit evidence. By offering both high-level visibility and granular evidence, CARA fosters trust among stakeholders, including regulators, customers, and internal compliance officers (Odeh et al., 2024).

5. Logic Synthesis and Assurance (CLSA & CARA)

The Compliance Logic Synthesis Agent (CLSA) and Continuous Assurance & Reporting Agent (CARA) constitute the operational core of the AC2C framework, wherein abstract regulatory interpretations and real-time data mappings coalesce into actionable, automated compliance mechanisms. In the HealthSphere case study, the CLSA demonstrates its capacity to synthesize precise, executable logic tailored to the organization's multi-regulatory environment, while CARA ensures ongoing verification and transparent reporting. This integration not only addresses the "regulation-to-code" bottleneck but also facilitates proactive risk mitigation in dynamic IoT ecosystems, as evidenced by the framework's handling of specific compliance scenarios across HIPAA, GDPR, and NIS2.

5.1 HIPAA

The CLSA synthesizes tests to verify that Protected Health Information (PHI) stored in the US is encrypted at rest and that access controls adhere to the HIPAA Security Rule. Drawing from the Compliance Knowledge Graph (CKG) generated by the RDA, which identifies HealthSphere as a "Business Associate" under HIPAA, the CLSA generates Python-based scripts or Rego policies that automatically check for encryption standards (e.g., AES-256) on cloud storage instances and enforce role-based access controls (RBAC) aligned with the minimum necessary principle (Said et al., 2024). In the case study, this logic was applied to HealthSphere's cloud analytics platform, where PHI from wearable devices is processed, ensuring that any unauthorized access attempts trigger immediate flags. This automation reduces the manual audit burden, which traditionally consumes significant IT resources (Ghafari et al., 2024), and aligns with HIPAA's emphasis on safeguarding electronic PHI against unauthorized disclosure.

5.2 GDPR

The CLSA develops tests to ensure the presence of a valid data transfer mechanism, such as Standard Contractual Clauses, for EU-to-US data transfers and to confirm that explicit consent has been obtained for processing special category data. Utilizing the DFIA's mapping of cross-border data flows, such as Flow B involving EU health data routed to US servers, the CLSA incorporates semantic models to generate compliance checks that validate transfer mechanisms in accordance with GDPR Articles 44-50. For example, it synthesizes code to examine metadata logs for SCC implementation and to verify consent records in the mobile application's database, ensuring they comply with the requirements for explicit, informed consent under Article 9 for sensitive health data. This process underscores the framework's efficacy in addressing regulatory ambiguities, such as "appropriate safeguards," by contextualizing them within HealthSphere's operational setup, thereby mitigating potential fines and facilitating real-time adjustments to data processing activities (Said et al., 2024).

5.3 NIS2

The CLSA integrates tests to effectively manage supply chain risks by ensuring that a third-party vendor has undergone assessment, the firmware's SBOM has been scanned for vulnerabilities, and the update is cryptographically signed. Guided by the DFIA's identification of Dependency C, the CLSA employs the RDA's CKG to align with NIS2's requirements for essential entities. This process generates OPA Rego policies that automate vendor risk assessments, SBOM vulnerability scans through integrated tools, and signature validation using cryptographic hashes (Chiara, 2022). In the HealthSphere scenario, this logic was applied to wearable device updates, ensuring compliance with NIS2's emphasis on cybersecurity resilience and supply chain security. By automating these checks, the framework mitigates latent risks that static models may overlook, thereby supporting continuous monitoring in response to evolving IoT threats (Kilaru et al., 2024).

5.4 Incident Response

In a simulated breach scenario, the CARA autonomously generates three distinct, parallel reports to comply with the varying notification deadlines stipulated by NIS2 (24-hour early warning), GDPR (72-hour notification), and HIPAA (60-day notification). Upon identifying a violation, such as a simulated unauthorized access to PHI within the cloud platform, CARA executes the CLSA-synthesized tests at a high frequency, tracing the incident back to specific

CKG nodes (e.g., breaches of the HIPAA Security Rule or obligations under GDPR Article 33) (Odeh et al., 2024). Subsequently, it produces immutable, time-stamped reports with cryptographic hash chains: a concise early warning for NIS2 regulators within 24 hours, a detailed GDPR notification including impact assessments within 72 hours, and a comprehensive HIPAA report for affected individuals and authorities within 60 days. These outputs encompass enriched alerts with event logs, regulatory citations, and remediation recommendations, thereby fostering explainability and accountability (Schmitz-Berndt, 2023). This capability underscores CARA's role in transforming reactive compliance into a proactive, auditable process, which is particularly vital for IoT systems where the velocity of incidents necessitates rapid response.

The performance of the CLSA and CARA in the HealthSphere case study substantiates the effectiveness of the AC2C framework in synthesizing and ensuring compliance logic across overlapping regulations. By automating the translation of regulatory intent into executable code and facilitating continuous verification, the framework not only addresses the compliance gap but also enhances organizational resilience (Odeh et al., 2024). This, in turn, reduces financial burdens and fosters innovation within regulated IoT environments.

Discussion

The HealthSphere case study provides a practical demonstration of the AC2C framework's capacity to navigate a complex, multi-jurisdictional regulatory landscape. The scenario, involving a global healthcare provider, effectively illustrates how the framework addresses the intricate compliance challenges posed by the simultaneous application of GDPR, HIPAA, and the NIS2 Directive.

The initial actions of the Regulatory Deconstruction Agent (RDA) are foundational to the framework's success. By correctly classifying HealthSphere as a "Business Associate" under HIPAA, a "Data Controller" under GDPR, and an "Important Entity" under NIS2, the RDA showcases a sophisticated comprehension of legal and organisational contexts (Sadri, 2024). This automated classification overcomes a significant manual hurdle in traditional compliance, ensuring that the subsequent analysis is grounded in an accurate understanding of the organisation's specific legal obligations (Sardana, 2024).

Furthermore, the IoT Data Flow Intelligence Agent (DFIA) provides the critical real-world context that static compliance approaches lack. The agent's mapping of key data flows reveals immediate and significant compliance risks (Pasquier et al., 2018). For instance, identifying the transfer of EU health data to the United States (Flow B) directly triggers the stringent cross-border data transfer requirements of GDPR (Mulder & Tudorica, 2019). Similarly, mapping the firmware update process from a third-party vendor in India (Dependency C) highlights a crucial supply chain vulnerability that falls squarely within the scope of NIS2 (Schip, 2024). This dynamic environmental awareness allows the framework to move beyond theoretical policy checks to address actual operational risks as they emerge.

The synergy between the RDA's legal intelligence and the DFIA's operational intelligence enables the AC2C framework to effectively bridge the regulation-to-code gap (Meroni et al., 2018). The framework synthesises these two streams of information to generate precise, automated compliance checks that are directly relevant to the organisation's activities. The case study demonstrates that the framework can translate the abstract principles of data protection,

security, and resilience into concrete, verifiable actions, thereby offering a viable path toward continuous, autonomous compliance in the modern IoT ecosystem.

References

- 1) Alles, M., Brennan, G., & Kogan, A. (2018). Continuous monitoring of business process controls: A pilot implementation of a continuous auditing system at Siemens. <https://www.emerald.com/insight/content/doi/10.1108/978-1-78743-413-420181010/full/html>
- 2) Barron, R., Eren, M., & Serafimova, O. (2025). Bridging Legal Knowledge and AI: Retrieval-Augmented Generation with Vector Stores, Knowledge Graphs, and Hierarchical Non-negative Matrix Factorization. <https://arxiv.org/abs/2502.20364>
- 3) Chaganti, K. C. (2025). Advancing ai-driven threat detection in iot ecosystems: Addressing scalability, resource constraints, and real-time adaptability. Institute of Electrical and Electronics Engineers (IEEE). <https://doi.org/10.36227/techrxiv.173738307.73168902/v1>
- 4) Chiara, P. G. (2022). The IOT and the new EU Cybersecurity Regulatory Landscape. *International Review of Law, Computers & Technology*, 36(2), 118–137. <https://doi.org/10.1080/13600869.2022.2060468>
- 5) Chiara, P., 2022. The Network and Information Systems Directive 2 (NIS2): Enhanced cybersecurity obligations for EU entities. *Journal of European Law and Policy*, 15(3), pp. 45-62.
- 6) Coulter, R., & Pan, L. (2018). Intelligent agents defending for an IoT world: A review. *Computers & Security*. <https://www.sciencedirect.com/science/article/pii/S0167404817302511>
- 7) Echenim, K. U., & Joshi, K. P. (2023). IoT-Reg: A comprehensive knowledge graph for real-time IoT data privacy compliance. 2023 IEEE International Conference on Big Data (BigData), 2897–2906. <https://doi.org/10.1109/bigdata59044.2023.10386545>
- 8) Ettaloui, N., Arezki, S., & Gadi, T. (2023). An Overview of Blockchain-Based Electronic Health Records and Compliance with GDPR and HIPAA. *Data and Metadata*. <https://www.semanticscholar.org/paper/292ed863c0463a9b230c966fd25f61b47f7d0717>
- 9) Ghafari, F., Shourangiz, E., & Wang, C. (2024). Cost effectiveness of the industrial internet of things adoption in the U.S. manufacturing smes. *Intelligent and Sustainable Manufacturing*, 1(1), 10008–10008. <https://doi.org/10.35534/ism.2024.10008>
- 10) Ghafari, S., Smith, J. and Patel, R., 2024. Cost analysis of compliance in IoT-driven enterprises. *International Journal of Cybersecurity Management*, 8(1), pp. 12-25.
- 11) Halgamuge, M. N., & Niyato, D. (2025). Adaptive Edge Security Framework for dynamic IOT security policies in diverse environments. *Computers & Security*, 148, 104128. <https://doi.org/10.1016/j.cose.2024.104128>
- 12) Hornos, M. J., & Quinde, M. (2024). Development methodologies for IOT-based systems: Challenges and Research Directions. *Journal of Reliable Intelligent Environments*, 10(3), 215–244. <https://doi.org/10.1007/s40860-024-00229-9>
- 13) Junaid, S. B., Imam, A. A., Balogun, A. O., De Silva, L. C., Surakat, Y. A., Kumar, G., Abdulkarim, M., Shuaibu, A. N., Garba, A., Sahalu, Y., Mohammed, A., Mohammed,

- T. Y., Abdulkadir, B. A., Abba, A. A., Kakumi, N. A., & Mahamad, S. (2022). Recent advancements in emerging technologies for Healthcare Management Systems: A survey. *Healthcare*, 10(10), 1940. <https://doi.org/10.3390/healthcare10101940>
- 14) Kilaru, S., Zhang, L. and Chen, H., 2024. Continuous discovery in IoT: Leveraging network traffic and SBOMs for security. *IEEE Internet of Things Journal*, 11(5), pp. 890-902.
- 15) Kilaru, M., Maheswari, P., Boddepalli, E., Venkataramana, K., Patel, J. D., & Sharma, M. K. (2024). IoT Services and Intelligence: Empowering the Internet of Things with Real-Time Data Analytics and decision-making. 2024 International Conference on Trends in Quantum Computing and Emerging Business Technologies, 1–5. <https://doi.org/10.1109/tqcebt59414.2024.10545282>
- 16) Kulkarni, V., Sunkle, S., Kholkar, D., Roychoudhury, S., Kumar, R., & Raghunandan, M. (2021). Toward automated regulatory compliance. *CSI Transactions on ICT*, 9(2), 95–104. <https://doi.org/10.1007/s40012-021-00329-4>
- 17) Li, J., Maiti, A., & Fei, J. (2023). Features and scope of regulatory technologies: Challenges and opportunities with industrial internet of things. *Future Internet*, 15(8), 256. <https://doi.org/10.3390/fi15080256>
- 18) Li, S., Chen, J., Yao, R., Hu, X., Zhou, P., Qiu, W., Zhang, S., Dong, C., Li, Z., Xie, Q., & Yuan, Z. (2025). Compliance-to-Code: Enhancing Financial Compliance Checking via Code Generation. *ArXiv*. <https://arxiv.org/abs/2505.19804>
- 19) Meroni, G., Baresi, L., Montali, M., & Plebani, P. (2018). Multi-party business process compliance monitoring through IoT-enabled artifacts. *Information Systems*. <https://www.sciencedirect.com/science/article/pii/S0306437917301242>
- 20) National Institute of Standards and Technology. (n.d.). AI risk management framework. Retrieved July 4, 2025, from <https://www.nist.gov/itl/ai-risk-management-framework> 897–2906. <https://doi.org/10.1109/bigdata59044.2023.10386545>
- 21) Odeh, A., Abu Taleb, A., Alhajajeh, T., Aparicio, F., Hamed, S., Al Daradkeh, N., & Ali Al-Jarallah, N. (2024). Data privacy and compliance in IoT. In *Advances in Information Security, Privacy, and Ethics* (pp. 128–144). IGI Global. <https://doi.org/10.4018/979-8-3693-3451-5.ch006>
- 22) Odeh, A., Farooqi, M. and Khan, S., 2024. Real-time compliance auditing in IoT: Challenges and solutions. *Journal of Cybersecurity Research*, 9(2), pp. 67-80.
- 23) Pasquier, T., Singh, J., Powles, J., Evers, D., Seltzer, M., & Bacon, J. (2017). Data provenance to audit compliance with privacy policy in the Internet of Things. *Personal and Ubiquitous Computing*, 22(2), 333–344. <https://doi.org/10.1007/s00779-017-1067-4>
- 24) Sadri, M. (2024). HIPAA: A Demand to Modernize Health Legislation. *The Undergraduate Law Review at UC San Diego*. <https://escholarship.org/uc/item/9gp2n52k>
- 25) Said, A., Mahmoud, K. and El-Sayed, M., 2024. HIPAA and GDPR compliance in healthcare IoT: A comparative analysis. *Health Informatics Journal*, 30(1), pp. 56-72.
- 26) Said, A., Yahyaoui, A., & Abdellatif, T. (2024). HIPAA and GDPR compliance in IOT healthcare systems. *Communications in Computer and Information Science*, 198–209. https://doi.org/10.1007/978-3-031-55729-3_16

- 27) Sardana, A., Sethuraman, S., & Kalyanasundaram, P. D. (2024). Compliance-as-code 2.0: Orchestrating regulatory operations with agentic AI. *Journal of Artificial Intelligence General Science (JAIGS)* ISSN:3006-4023, 5(1), 546–563.
<https://doi.org/10.60087/jaigs.v5i1.366>
- 28) Sardana, J. (2024). Automating Global Trade Compliance through Product Classification Systems. <https://inlibrary.uz/index.php/tajmei/article/view/78575>
- 29) Schip, M. van 't. (2024). The Regulation of Supply Chain Cybersecurity in the NIS2 Directive in the Context of the Internet of Things. *SSRN Electronic Journal*.
<https://doi.org/10.2139/ssrn.4848048>
- 30) Schmitz-Berndt, S., 2023. Timely incident response in multi-regulatory environments: Lessons from GDPR and NIS2. *European Journal of Information Systems*, 32(4), pp. 201-215.
- 31) Shahin, M., Hosseinzadeh, A., & Chen, F. F. (2025). A two-stage hybrid federated learning framework for privacy-preserving IoT anomaly detection and classification. *IoT*, 6(3), 48. <https://doi.org/10.3390/iot6030048>
- 32) Shen, T., Zhang, F., & Cheng, J. (2022). A comprehensive overview of knowledge graph completion. *Knowl. Based Syst.*
<https://linkinghub.elsevier.com/retrieve/pii/S095070512200805X>
- 33) Voss, W. (2019). Cross-border data flows, the GDPR, and data governance. *Wash. Int'l LJ*. https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/pacrimlp29§ion=23
- 34) Schmitz-Berndt, S. (2023). Defining the reporting threshold for a cybersecurity incident under the NIS directive and the NIS 2 directive. *Journal of Cybersecurity*, 9(1). <https://doi.org/10.1093/cybsec/tyad009>
- 35) Singh, K., & Singh, B. (2024). Multimodal Data Retrieval Challenges and their Countermeasures Using Novel Integrated Data Mining and Fusion System (IDMFS). In *Emerging Trends in IoT and Computing Technologies* (pp. 286–292). CRC Press.
<https://doi.org/10.1201/9781003535423-48>
- 36) Xu, L., Lu, L., Liu, M., Song, C., & Wu, L. (2024). Nanjing Yunjin intelligent question-answering system based on knowledge graphs and retrieval augmented generation technology. *Heritage Science*, 12(1). <https://doi.org/10.1186/s40494-024-01231-3>
- 37) Zaman, S. A. (2023, August). (PDF) internet of things (IOT) data protection and security concerns -Review. *ResearchGate*. 10.13140/RG.2.2.12361.52321