Architecting Regulatory Agility: The GCOA Framework for Decoupling Compliance Logic in Global E-commerce Systems.

Sureshkumar Karuppuchamy
eBay Inc.
San Jose, CA, USA
Email: skaruppuchamy05@gmail.com

Abstract*—Global online marketplaces face unprecedented pressure to embed evolving regulations into their software architecture. This paper presents an architectural case study of a large e-commerce platform's response to the EU General Product Safety Regulation (GPSR) and the Digital Services Act (DSA) by treating legal mandates as explicit design constraints. We propose the Generalized Compliance Orchestration Architecture (GCOA), a reusable, two-layer framework that achieves regulatory agility by decoupling compliance logic from core business functions. The architecture introduces two key, modular components: a configurable Regulatory Policy Engine (RPE) and a State Machine Manager (SMM). The Preventive Compliance Layer enforces ex ante product safety rules (GPSR) via a multi-stage validation pipeline at listing creation, while the Reactive Governance Layer manages ex post content governance (DSA) using automated notice-and-action workflows. Key technical elements include the SMM-driven "Listings On Hold" lifecycle state machine with an integrated appeal process, prioritized event handling for trusted flaggers, and an immutable Audit & Transparency Ledger (ATL) for verifiable accountability. The platform achieved on-time compliance by the regulatory deadlines with minimal impact on user experience. This work demonstrates that regulatory requirements can be met at scale through modular, automated architecture. *This work is presented as an architectural case study rather than a controlled experimental evaluation; its contribution lies in demonstrating how regulatory requirements can be operationalized as modular, scalable system components at production scale.

***Index Terms*—regulatory compliance; systems architecture; e-commerce; compliance-as-code; audit logging; state machines; Digital Services Act (DSA); General Product Safety Regulation (GPSR)**

I. Introduction and Regulatory Context

In 2024–2025, a global e-commerce marketplace faced hard EU deadlines for product safety and online content governance [1][2]. The General Product Safety Regulation (GPSR) (effective December 2024) requires marketplaces to ensure non-food consumer products include verified safety information and to remove unsafe items swiftly[3]. The Digital Services Act (DSA) (effective February 2024) mandates notice-and-action, user statements of reasons, and an appeal mechanism [2][4], with penalties up to 6% of global revenue for non-compliance [5]. Rather than treat GPSR/DSA as isolated policy changes, the platform used them as first-class architectural drivers, translating legal duties into design constraints. This produced the Generalized Compliance Orchestration Architecture (GCOA), embedding preventive checks and reactive governance into core services and enabling on-time compliance while establishing a template for future regulatory agility [6]. While this paper focuses on a large-scale marketplace deployment, the architectural mechanisms described—policy externalization, state-machine–driven governance, and append-only auditability—are not inherently scale-dependent. Section III discusses constraints and adaptations required for smaller or resource-constrained platforms.

Unlike prior policy-driven or logging-centric approaches, GCOA integrates preventive enforcement, reactive governance, and legally mandated due-process workflows within a single stateful compliance orchestration layer.

Although motivated by EU regulations, GCOA addresses regulatory compliance as a recurring architectural problem rather than a jurisdiction-specific solution.

II. Architecture

A. GCOA Architecture Overview

The Generalized Compliance Orchestration Architecture (GCOA) comprises two integrated layers that modularize regulatory enforcement. The Preventive Compliance Layer implements proactive controls to block non-compliant listings before publication, fulfilling the GPSR's ex ante safety requirements. The Reactive Governance Layer manages content already live, fulfilling the DSA's notice-and-action obligations. Together, the layers enforce compliance across the full listing lifecycle. At the core are four key components:

Regulatory Policy Engine (RPE): A configurable rules engine that validates listings against compliance policies. It supports synchronous blocking checks at creation and asynchronous background scans. New or updated rules are deployed through configuration, not code, enabling rapid adaptation to evolving regulations[8].

State Machine Manager (SMM): A workflow orchestrator governing allowed listing states (e.g., Active → On Hold → Pending Appeal → Reinstated/Removed). When the RPE flags a violation, it triggers a state change via the SMM, enforcing a formal lifecycle model for governance events.

Compliance Service (C-Service): A dedicated microservice encapsulating the RPE and SMM. It mediates between the core

listing platform and dependent systems (databases, notifications, etc.), decoupling compliance logic from business code. This "compliance-as-code" design improves maintainability—legal updates typically require only configuration or C-Service changes.

Audit & Transparency Ledger (ATL): An append-only, cryptographically linked ledger recording all compliance-related actions—rule outcomes, notices, state changes, and appeals. The ATL enables automated transparency reporting and tamper-evident auditability, directly supporting DSA accountability requirements[9].

B. Technical Specificity: RPE Technologies and ATL Cryptography

Regulatory Policy Engine (RPE): The RPE is an event-driven rules engine integrated into the platform's microservices. It executes both synchronous validations (at listing creation) and asynchronous policies (background checks). A configurable rules framework—such as Drools BRMS or Open Policy Agent—allows new rules to be added via configuration rather than code. The engine subscribes to key platform events: for example, a listing created event triggers inline policy checks, while update or periodic events launch background re-scans. An event-driven architecture using background workers and message queues enables high-volume rule evaluations in near real time without blocking user flows. Critical GPSR checks run synchronously at submission, while resource-intensive tasks (e.g., image analysis, recall database lookups) run asynchronously. This design balances compliance assurance and seller experience. The stack uses standard, proven components—a scalable rules engine for decision logic and a pub/sub or streaming system (e.g., Kafka, AWS SNS/SQS)—ensuring reliability and rapid adaptation as regulations evolve[7].

Audit & Transparency Ledger (ATL): The ATL is an immutable, append-only ledger that records every compliance event with cryptographic integrity. Each entry stores a timestamp, event data, and a hash pointer to the previous record, forming a tamper-evident chain using a secure hash algorithm (e.g., SHA-256). In summary, the ATL provides an append-only, cryptographically verifiable audit trail suitable for regulatory accountability and transparency reporting.

C. Preventive Compliance Layer (GPSR)

The first layer of GCOA enforces preventive compliance, ensuring product safety and regulatory completeness before a listing goes live. To meet GPSR obligations, the platform embedded structured safety and compliance checks directly into the listing creation workflow—transforming it into a gatekeeper that blocks non-compliant products from reaching buyers.

To meet GPSR requirements, the listing data model was extended with mandatory structured fields covering manufacturer identity, standardized product identifiers, localized safety warnings, and required compliance documentation. Listings missing these fields are rejected at creation time.

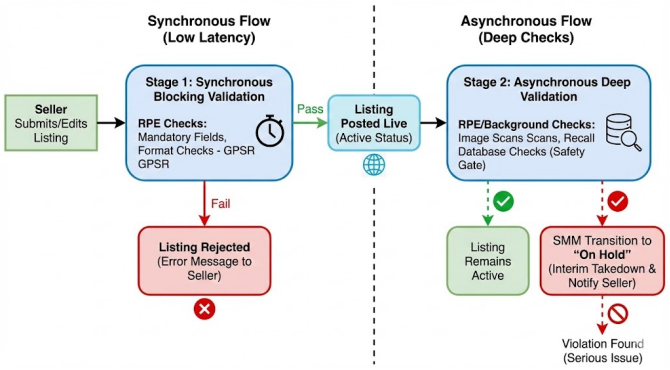

Fig. 1. Two-stage listing validation: Stage 1 (synchronous) blocks missing/invalid GPSR fields at submission; Stage 2 (asynchronous) runs post-publish (image, recall, and text scans) and can move listings to On Hold.

1) Two-Stage Validation Pipeline: Adding new data fields alone was insufficient—the platform also had to validate their accuracy and completeness without degrading the seller experience. Heavy synchronous checks could slow submissions, so the GCOA introduced a two-stage validation pipeline balancing compliance coverage with performance.

Synchronous Blocking Validation:

When a seller submits a new or edited listing, the system performs fast, deterministic "blocker" checks within the main request flow. Executed by the RPE, these validations ensure all mandatory GPSR fields are present and well-formed—for example, confirming an EU-based responsible address, valid identifiers, and attached documentation. Listings failing these checks are immediately rejected with corrective prompts. This inline gatekeeping adds only about 100 ms of latency; median creation time remained under two seconds—an acceptable trade-off for guaranteed baseline compliance. By the GPSR effective date, no listing lacking essential safety data could go live.

Asynchronous Deep Validation:

After publication, the Compliance Service launches deeper background scans that run independently of the user transaction. These asynchronous checks include: Content and image analysis: verifying safety markings, age labels, and detecting missing warnings through image recognition. Text analysis: using NLP to confirm that safety warnings and instructions appear in all required languages and formats. External data integration: cross-referencing product identifiers against recall databases such as the EU's Safety Gate alerts for dangerous goods.

If a severe issue emerges—e.g., a recall match or missing label—the listing is automatically transitioned by the SMM into an On Hold state, removing it from buyer view pending review.

Although this design introduces a short delay before removal, detections typically occur within seconds, maintaining both compliance and a smooth seller experience. The approach intentionally trades strict immediacy for throughput and usability, while still satisfying the DSA/GPSR principle of acting

"without undue delay."

D. Reactive Governance Layer (DSA)

The second GCOA layer handles moderation and due process *after* a listing is live. While the preventive layer blocks unsafe listings up front, the reactive layer addresses items that become suspect or illegal post-publication (e.g., user reports, authority notices). To meet DSA requirements for swift, transparent action, we implemented a stateful workflow centered on an intermediate *On Hold* state, with services for notice intake, appeals, notifications, and audit logging.

1) "Listings On Hold" Lifecycle State Machine: The *On Hold* state, managed by the SMM, quarantines listings under investigation. Upon a credible notice or automated policy flag, the listing transitions to *On Hold*: it is immediately removed from buyer access (hidden from search and purchase) while all data is preserved for review. This satisfies the DSA's "disable access without undue delay" mandate while retaining evidence for due process.

We formalized a finite-state machine (FSM) with the following states (Figure 2):

Active — Default, live state after GPSR checks; assumed compliant.

On Hold — Entered on credible notice or automated flag. Buyer access is disabled immediately; the system logs the event to the ATL and sends the seller a *statement of reasons*. Data remains intact but inaccessible.

Pending Appeal — Set when a seller contests the action. The listing remains off-market while Trust & Safety reviews the case. The ATL records appeal receipt and review status.

Reinstated — If review finds no violation, the listing returns to *Active*. Notifications are sent to the seller (and, optionally, the reporter); outcomes are logged in the ATL.

Removed — If the violation is confirmed or no appeal is filed within the defined window, the listing is permanently removed from buyer access. The seller receives a final decision notice; the ATL records the resolution. Repeat removals can trigger account-level enforcement (repeat-offender handling).

This state machine operationalizes DSA notice-and-action with clear transitions, auditable artifacts (ATL), and a built-in path for timely seller appeals, ensuring both rapid mitigation and procedural fairness.

2) Automated Notice Handling and Trusted Flaggers: To operationalize the reactive workflow, the platform implemented an automated notice-intake system within the Compliance Service. All reports—whether from regular users (via Report Listing), government agencies, consumer organizations, or internal automated detectors—flow into a unified queue for evaluation. The RPE processes each notice to determine the appropriate response: low-confidence reports are routed for manual review, while high-confidence or policy-matched cases trigger immediate On Hold transitions.

A special DSA category, Trusted Flaggers (Art. 22), receive prioritized handling. The system maintains a registry of accredited entities; any notice submitted by these sources is automatically marked high priority and processed without delay.

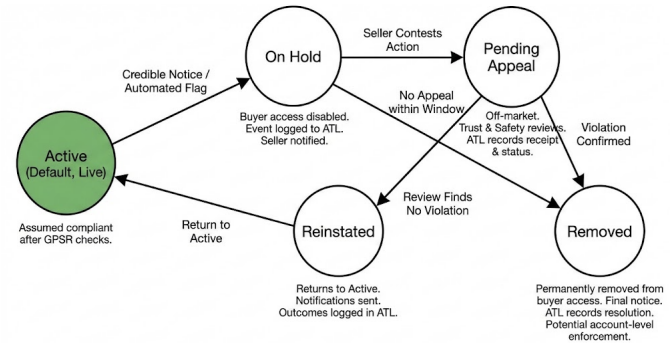

Fig. 2. DSA "Listings On Hold" State Machine

In practice, the RPE rule If Notice.Sender = TrustedFlagger $\rightarrow$ Set Listing Status = On Hold ensures near-instant action. This yielded a median time-to-action < 12 hours (often minutes for critical flags) versus $\approx$36 hours for general user reports. All takedowns, regardless of source, still generate seller notifications and allow appeals—trusted does not mean unquestioned, only prioritized.

To prevent misuse of the notice system, the platform added rate-limits and reporter reputation tracking. Users generating repeated invalid reports (no-action outcomes) are automatically de-prioritized or temporarily suspended from submitting notices. These governance rules, encoded in the RPE, ensure consistent enforcement, transparency, and speed—no report is lost, and every decision is logged and auditable through the ATL.

The Compliance Service's RPE triages all notices. Trusted Flagger reports are high priority and immediately move listings On Hold. General user reports go through validation; credible ones trigger On Hold, others queue for manual review. This fast-track/triage design scaled to thousands of reports per day with automated rules and targeted escalations, and all outcomes are logged in the ATL for transparency and audit.

Sellers are notified of On Hold actions with a statement of reasons and may appeal via a self-service interface. Appeals are reviewed by a human moderator; outcomes are logged in the ATL. Average appeal resolution time was $\approx$3.2 days, with a 15% reinstatement rate.

3) Immutable Audit & Transparency Ledger (ATL): A cornerstone of GCOA's reactive layer is the Audit & Transparency Ledger (ATL) underpins the trust and accountability of the system. The ATL is essentially a secure, tamper-evident log of all actions taken under the compliance workflows. Every time a listing's state changes, or a notice is received, or a notification is sent, an entry is appended to the ledger. Each log entry includes a timestamp, the event details (e.g. "Listing 12345 moved to On Hold due to report by user X"), and is cryptographically linked to the previous entry via a hash, forming a chain.

In practice, we implemented the ATL as a dedicated database table with an append-only policy and used a hashing scheme to link entries. We did not deploy a full decentralized blockchain (not needed for an internal system), but the concept is similar to using blockchain for compliance evidence as suggested by

industry research. The ledger can be exported or audited by external parties if required – for example, if a regulator asks to verify how a particular notice was handled, we can provide the chain of events from notice receipt to takedown time to final resolution. This proves that we acted "without undue delay" and followed the correct process, as every step's timestamp is recorded.

Importantly, the ATL also streamlines compliance reporting. The DSA requires annual transparency reports (Article 15) with statistics on content moderation (e.g., number of notices handled, median time to action, number of removals and reinstatements, etc.). With the ATL, generating these metrics became trivial – we could query the ledger for the relevant data instead of relying on disparate logs. For instance, in our 2025 report we could state: "95,245 notices were processed, median takedown time was 16 hours, and 12% of seller appeals resulted in content reinstatement," all computed directly from ledger entries. This level of transparency not only satisfies regulators but also provides internal insight. Teams have started analyzing ATL data to identify patterns (e.g., which rules trigger the most appeals, which product categories see frequent violations) to continuously refine our policies.

From a performance standpoint, writing to the ATL is lightweight – just a single database insert and hash computation per event – so it did not noticeably impact system throughput. We implemented batch signing of ledger entries periodically to further harden integrity (and are considering external timestamping services for even stronger proofs). Overall, the ATL gave us measurable, auditable compliance. It turned what could have been a black-box moderation process into a verifiable chain of decisions, reinforcing accountability.

E. Architectural Trade-offs and Results

Designing and implementing the GCOA required navigating classic engineering trade-offs to satisfy the stringent legal requirements without degrading the platform's performance or user satisfaction. We highlight a few key considerations and outcomes:

Latency vs. Thoroughness:Adding compliance checks inevitably introduces some overhead. We chose to enforce critical GPSR checks synchronously (to block non-compliant listings immediately) and accept a slight increase in listing submission time (~100ms per listing on average). This was deemed a worthwhile price for ensuring no unsafe product is listed for sale. By offloading heavier checks to asynchronous processes, we kept the seller experience fast. For reactive takedowns, we favored an event-driven asynchronous model (using background workers and message queues) to process large volumes of removal events without slowing down the site. This can introduce a short propagation delay (a listing might remain visible for a few seconds while the "On Hold" event travels through the system), but we measured that 99% of removals propagated within ~30 seconds, which is operationally "without undue delay". We explicitly traded off strong immediate consistency in favor of scalability and resilience, guided by the insight that eventual consistency is acceptable for compliance actions as long as delays are kept very short.

Modularity vs. Coupling:A major goal was to decouple the regulatory logic from core marketplace logic. The introduction of the Compliance Service (with RPE and SMM) achieved high modularity – the listing service doesn't need to know the details of GPSR or DSA rules, it just calls the compliance API and reacts to its responses. The trade-off is a bit more inter-service communication and potential consistency lag (e.g., a slight delay for the listing service to learn a listing was put On Hold). We mitigated this by caching compliance status in read operations to avoid showing buyers an On-Hold listing in that brief window. The benefits of modularity have been clear: when new rules or adjustments come (as happened with a mid-2025 regulatory update requiring additional seller disclosures), we were able to implement most changes within the Compliance Service and simply update configuration in the RPE, with minimal changes to other services. This confirms that separating concerns (legal logic vs. business logic) increases agility, at the cost of some added complexity in integration.

Performance and Scalability:We rigorously tested the system to ensure it could handle the scale. The platform sees tens of thousands of listings created per minute at peak; the synchronous validation stage was load-tested and scaled horizontally to ensure it could keep up without creating a bottleneck. For content removals, we simulated scenarios like a mass recall of 100,000 listings in a short period – the event-driven architecture processed them without significant impact on the live site's performance, whereas a naive synchronous removal approach could have caused noticeable slowdowns. We also treated the Compliance Service as a critical, highly-available component – with redundancy and a "fail-safe" mode where if the service is down, the system errs on the side of blocking actions (fail closed) to avoid letting potentially illegal content slip through. This emphasis on reliability ensures that compliance features have uptime on par with other mission-critical systems (like payments).

Accuracy vs. Overreach:Automating compliance means finding the right balance between catching all violations and minimizing false positives that burden users. Our initial deployment showed a false positive rate (listings wrongly taken down and later reinstated) of about 15% for automated flags. While not trivial, this rate was manageable and in line with expectations for a first iteration of rules. Continuous improvements, such as refining RPE rules and incorporating feedback from overturned appeals, are underway to reduce false positives. On the flip side, our "recall" of true violations (finding illegal content that would have previously gone unnoticed) jumped to ~90% with GCOA, from an estimated 50% baseline when we relied solely on user reports. This indicates a significant boost in catching bad content before it causes harm. We consider this an ongoing optimization problem: using more advanced AI for detection carefully, and possibly adaptive algorithms (as discussed in future work) to improve over time.

Overall, the results from the first few months of GCOA in production were very positive. The platform achieved

virtually 100% compliance with the GPSR by the deadline – essentially no new EU-bound listings were missing the required safety info. For the DSA, all required processes (notice handling, user notifications, transparency reporting) were in place by the enforcement date, and internal audits as well as external regulatory checks confirmed the platform's compliance. Crucially, these compliance features did not wreak havoc on the business metrics: listing creation rates, sales volume, and user engagement held steady, with only a tiny fraction of listings ever being interrupted by the new measures (at any given time, ~0.05% of active listings were On Hold). Seller feedback, initially cautious, turned largely positive when they saw that these changes led to a safer marketplace and were applied fairly (in a 2025 survey, over 80% of sellers understood the reasons for the new requirements and felt it had no negative impact or even improved buyer trust). This underscores that compliance can be an opportunity to strengthen a platform's value proposition (e.g., "we are a safer, more trustworthy marketplace") rather than just a regulatory burden.

F. Operational Evaluation and Reproducibility Scope

Although this work does not present a synthetic benchmark or open testbed, the architecture was evaluated under real production workloads. Key evaluation dimensions included: (i) synchronous listing latency overhead (median ≈ 100 ms), (ii) notice-to-action time under burst conditions (P99 < 30 seconds for automated actions), (iii) appeal resolution time (mean ≈ 3.2 days), and (iv) false-positive rate for automated enforcement ($\approx 15\%$ at initial deployment).

The event-driven design, rule configurations, and state-machine definitions are reproducible in principle using standard components (e.g., Kafka-class messaging, OPA/Drools-style rule engines, append-only ledgers). However, the datasets and traffic traces used in this study are proprietary and cannot be released; thus, this paper emphasizes architectural reproducibility over dataset replication.

G. Generality and Deployment Constraints

The GCOA pattern is most naturally suited to platforms with high regulatory exposure and continuous policy change. Smaller platforms may adopt a simplified variant by collapsing the Compliance Service into a modular library, limiting the number of lifecycle states, or using conventional append-only logs instead of cryptographically chained ledgers. The core principle—treating compliance logic as a separately evolvable subsystem—remains applicable even when individual components are simplified.

III. Architectural Decision Rationale and Trade-off Analysis

Designing the Generalized Compliance Orchestration Architecture (GCOA) required a series of deliberate architectural decisions to balance performance, modifiability, scalability, and regulatory accountability. This section summarizes the rationale behind key design choices and analyzes their trade-offs.

A. Decision Context

The GCOA was designed to meet stringent regulatory requirements (EU GPSR, DSA) under tight time constraints while maintaining seller experience and system performance. Core architectural forces included:

- **Agility:** Ability to add or update compliance rules rapidly without code changes.
- **Scalability:** Processing billions of listings and millions of policy events daily.
- **Auditability:** Providing tamper-evident evidence trails for regulators.
- **Modularity:** Minimizing coupling with existing listing and catalog systems.

B. Key Decisions and Alternatives

Table I summarizes selected design decisions, alternatives considered, and their trade-offs.

TABLE I
Architectural Design Trade-offs

Decision	Alternatives	Rationale (Chosen)	Trade-off
Centralized Compliance Service	Embed rules in each service	Enables isolated updates and reuse across multiple domains	Adds network hops, requires strict SLAs between services
Two-stage Validation Pipeline	All-synchronous checks	Preserves low seller latency while ensuring eventual completeness	Introduces short window before full validation
Immutable Audit Ledger (ATL)	Standard DB logging	Provides verifiable, append-only evidence for audit and reporting	Higher storage footprint and signing overhead
Configurable RPE (Drools/OPA)	Hard-coded rule logic	Enables compliance-as-code and rapid rule deployment	Slightly higher runtime cost for rule interpretation
Event-driven Messaging (Kafka/SNS)	Direct synchronous calls	Ensures elasticity and backpressure handling	Adds operational complexity (queue management)

C. Impact on Quality Attributes

Each decision was evaluated against key quality attributes:

- **Performance:** End-to-end listing latency increased by only ≈0.1 s on average while scaling horizontally.
- **Modifiability:** New regulatory rules deployed via configuration without code redeployments.
- **Reliability:** Fail-safe design (fail-closed mode) prevented violations during service outages.
- **Auditability:** Cryptographic hash chains and batch signatures in the ATL ensured data integrity and external verifiability.

This structured rationale shows that the architecture balanced regulatory correctness with platform-scale operational demands.

D. Variability and Extensibility Mechanisms

The GCOA was designed as a configurable, reusable architecture capable of adapting to new regulatory domains.

Variation Points:

- The **Regulatory Policy Engine (RPE)** allows new policies to be introduced through configuration files or rule modules, decoupled from core services.
- The **State Machine Manager (SMM)** supports extensible lifecycle definitions—new states or transitions (e.g., "Under Review") can be added without redeploying the system.
- The **ATL schema** supports extension of event types for future laws such as sustainability or AI transparency requirements.

These mechanisms demonstrate that GCOA behaves as a product-line reference architecture where compliance domains are realized as configuration variants.

IV. Related Work and Positioning

Prior research on policy-driven and compliance-oriented architectures has explored decoupling regulatory rules from application logic. Elgammal et al. proposed compliance-by-design patterns for business processes, emphasizing rule externalization. Industry systems such as Oracle GRC and SAP Compliance Frameworks focus on after-the-fact reporting, while GCOA integrates compliance enforcement directly into runtime workflows.

Policy decision architectures like Google's Zanzibar demonstrate the scalability of centralized rule evaluation for authorization; GCOA applies a similar separation-of-concerns principle to regulatory compliance. Emerging "compliance-as-code" approaches (e.g., OPA, NIST OSCAL) share the same philosophy of declarative, auditable policy enforcement. GCOA extends these ideas with a two-layer architecture that unifies preventive (GPSR) and reactive (DSA) compliance and embeds a cryptographically verifiable ledger for transparency.

This synthesis positions GCOA as a novel application of software architecture principles—layering, modularity, and auditability—to regulatory engineering at internet scale.

Unlike general-purpose authorization systems such as Zanzibar or policy frameworks like OPA, GCOA integrates policy evaluation with explicit lifecycle governance and legally mandated due-process workflows. Similarly, while Certificate Transparency–style ledgers inspire the ATL's append-only design, GCOA applies these concepts to regulatory accountability rather than security certificates. The novelty of GCOA lies not in any single mechanism, but in their coordinated application to end-to-end regulatory enforcement across preventive and reactive domains.

Limitations

This paper does not propose new machine-learning models, formal verification techniques, or optimized detection algorithms. Instead, it assumes the existence of external signal providers (rules, classifiers, human reviewers) and focuses on how their outputs are governed, audited, and operationalized at scale. Future work may explore tighter integration with learning-based decision systems.

V. Conclusion

Treating regulation as architecture proved to be a powerful strategy for our global marketplace. By elevating GPSR and DSA requirements to first-class design constraints, we achieved compliance on time while also improving the platform's resilience and user trust. The Generalized Compliance Orchestration Architecture (GCOA)we developed – with its policy engine, stateful moderation workflow, and immutable audit logging – offers a reusable template for implementing new regulations as modular services in large-scale software platforms.

Our experience shows that regulatory compliance can be transformed from a reactive scramble into a proactive engineering advantage. The GCOA has become a core part of our platform, essentially a "compliance middleware" layer that can be extended and tuned as laws evolve. This positions us to respond with agility to the next wave of digital regulations (whether in AI, privacy, or other domains) without reinventing the wheel each time. Ultimately, embracing compliance requirements as a design challenge has not only kept us out of legal trouble – it has driven architectural innovation and strengthened the trust in our marketplace ecosystem.

Author Contributions and AI Assistance Disclosure

This paper was drafted by the author with selective assistance from AI-based language models (e.g., OpenAI ChatGPT) to generate preliminary outlines, refine technical phrasing, and ensure stylistic consistency. All substantive ideas, design descriptions, experimental results, and interpretations were authored and validated by the human authors. The authors take full responsibility for verifying the accuracy, originality, and integrity of all content.

References

[1] European Parliament and Council of the European Union, "Regulation (EU) 2023/988 on general product safety (GPSR)," *Official Journal of the European Union*, L 135, pp. 1–51, May 23, 2023.

[2] European Parliament and Council of the European Union, "Regulation (EU) 2022/2065 on a Single Market for Digital Services (Digital Services Act)," *Official Journal of the European Union*, L 277, pp. 1–102, Oct. 27, 2022 (applies from Feb. 17, 2024).

[3] European Commission, "Safety Gate: the EU rapid alert system for dangerous non-food products," 2025.

[4] Freshfields Bruckhaus Deringer, "The new General Product Safety Regulation: A refresh of the EU's product safety framework," Dec. 6, 2024.

[5] Crowell & Moring LLP, "Notice and Action Mechanisms in the DSA — Balancing the Removal of Illegal Content and the Freedom of Expression," Client Alert, Feb. 16, 2024. [Online]. Accessed: Nov. 14, 2025.

[6] European Commission, "AI Act — Shaping Europe's digital future," Policy page (Regulation (EU) 2024/1689), 2025.

[7] J. Kreps, N. Narkhede, and J. Rao, "Kafka: A Distributed Messaging System for Log Processing," in *Proc. NetDB Workshop*, 2011.

[8] R. Pang *et al.*, "Zanzibar: Google's Consistent, Global Authorization System," in *USENIX ATC*, 2019.

[9] B. Laurie, E. Messeri, and R. Stradling, "Certificate Transparency Version 2.0," RFC 9162, IETF, Dec. 2021. (Obsoletes RFC 6962.)