Federated Compliance Intelligence (FCI) for Enterprise Ecosystems

Luca Bianchi
Asst Professor

Abstract

Enterprise service ecosystems have emerged as a promising organizational model that enables companies to deliver integrated products and services by leveraging the capabilities of partners, customers, and third parties. However, the operation of enterprise-grade service ecosystems at scale is hampered by a number of technical and organizational barriers that create operational frictions. These barriers include the lack of trust between ecosystem partners, the existence of data and system silos, a patchwork of divergent policies and regulations, high costs related to compliance, and the presence of heterogeneous and often incompatible data standards.

Emerging engineering methods and standards are already addressing a subset of these barriers and can help support federation at the data level, control level, and decision level. A federated perspective enables the specification of protocols and procedures that facilitate compliance-related activities, such as privacy-preserving telemetry and incident response, and allows for the parallel execution of those activities across jurisdictional borders. Future research agendas and practical roadmaps can be aligned with these developments, with specific directions identified for the short term, medium term, and long term.

Keywords : Federated compliance, interoperability, data lineage, governance, scalability, policy engines, risk scoring, privacy-preserving computation.

1. Introduction

Enterprise service ecosystems—interconnected sets of organizations and service-provisioning platforms—are gaining traction as a scalable model for complex applications. Enterprises are outsourcing different components of their services to specialist service providers, also referred to as service dependencies, often hosted on public clouds. Organizations both in the supply and demand chains are facing increasingly stringent privacy and security regulations that are, at least on the face of it, incompatible with public-cloud deployments. Enterprises typically do not sell cloud-resident data to others, but rather select cloud services and service dependencies based on their risk exposure and appetite.

Compliance and regulatory requirements are evolving in the direction of data-residency constraints, cross-border data-transfer prohibitions, and domain-specific rules (e.g., healthcare, finance, children's sector). The enterprise ecosystems risk-compliance matrix is, however, not fully populated. Some risks—inherent to an enterprise ecosystem—are hard to minimize but easy to ignore by students of risk management; for others, these students lack the appropriate scientific models. These emergent properties of enterprises that are servicing their customers—much like the converging pattern of flocks of birds—appears intriguing. Achieving compliance in a federated context is a challenging problem, combining data engineering, security, and trust.Current gaps reside in the definition of adequate telemetries and the creation of dashboards that improve general-day-to-day operations. Compliance silos that only focus on data protection risk operations becoming painfully inefficient and extremely costly.

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

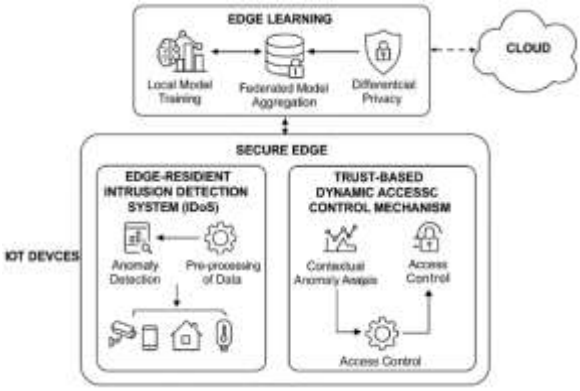

Fig 1: Federated Compliance Intelligence

1.1. Technical and Organizational Barriers

The success of enterprise service ecosystems depends on the capabilities and trustworthiness of constituent services, which in turn raises multiple technical, organizational, and economic challenges to scaling. Chief among these are the technical frictions that lead to the BUCKETS model. At its core, compliance is a major challenge that is hampered by operational inertia and escalating costs.

Data residency laws and other regulations act as a strong fricative against cross-border data flows, yet compliance checks are still often performed reactively and using the data itself, in spite of their shared regional context. Consequently, organizations in many regions and industries are investing in compliance controls. These costs disproportionately affect smaller organizations that are already disadvantaged when competing against larger actors, who may be able to absorb these costs or shift them elsewhere. Approaches that address the underlying technical and organizational barriers that create these frictions are therefore highly desirable.

Table 1: Technical and Organizational Barriers in Federated Compliance Ecosystems

Barrier Category	Description	Impact on Enterprise Ecosystems	Proposed Mitigation
Data Silos	Isolated organizational data repositories	Reduced interoperability and visibility	Federated data-sharing frameworks
Regulatory Fragmentation	Different regulations across jurisdictions	Increased compliance complexity	Cross-domain policy harmonization
Lack of Trust	Limited trust among ecosystem participants	Restricted collaboration and data exchange	Zero-trust governance and auditability
Heterogeneous Standards	Incompatible data and policy standards	Integration difficulties	Standardized interoperability protocols
High Compliance Cost	Expensive manual compliance operations	Reduced scalability for SMEs	Automated policy engines and telemetry
Cross-Border Restrictions	Data residency and transfer limitations	Operational delays and legal risks	Privacy-preserving federated architectures

1.2. Emerging Standards and Interoperability

Relevant standards, protocols, and policy models are maturing and should be leveraged to facilitate federated compliance. Strong support in the academic literature calls for a federated approach. Nevertheless, actual implementation faces frictions in achieving true federation between domains. Standards specific to compliance, or those user-declared to be relevant for compliance, can help. A thorough analysis of

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

existing data formats, exchange protocols, policy models, and standardization initiatives shows both their promise and the challenges encountered in adoption. Conformance tests may help overcome some of them. Proposed conformance tests complement operational technology for compliance but emphasize cooperation and incentive building for adoption by using the Department of Defense Security Automation Infrastructure Data Model as a lingua franca. A compromise of embracing prescribed structure without stifling flexible needs is desirable.

Interoperability brings together governance and compliance—the two fundamental aspects of enterprise risk management. A set of governance controls can support systemic compliance across domains and provide distance from compliance silos. There exist various governance-oriented standards for data that implement adequate mechanisms for inter-domain coordination. Support for compliance-related questions is intrinsic to these controls but often overlooked.

Equation 1: Federated Compliance Risk Score Aggregation

The paper discusses cross-domain risk scoring using normalized scores from multiple domains.

$$R_{global} = \sum_{i=1}^{n} w_i \cdot r_i$$

Where:

- R_{global}= aggregated federated compliance risk
- r_i= normalized risk score from domain i
- w_i= weighting factor for domain importance
- n= number of participating domains

$$R_{global} = \sum_{i=1}^{n} w_i \cdot r_i$$

1.3. Research Agendas and Practical Roadmaps

Three research agendas are proposed for the development, deployment, and operation of the required capabilities. The short-term agenda targets the realization of minimum viable products for proof-of-concept federated compliance applications by means of pilot implementations in enterprise service ecosystems. A mid-term agenda identifies open research challenges that must be addressed to enable broader adoption in complex service ecosystems. Finally, long-term strategic goals outline directions for federated compliance intelligence beyond compliance.

Initially, organizations seek to demonstrate that federated compliance approaches can deliver value within their specific environments by exposing them to tangible compliance-related challenges. Pilot implementations form dedicated usage environments for testing the integration of federated compliance capabilities into operational enterprise systems, determining resource requirements, and validating relative performance in real-world conditions. Ultimately, such experimentation should lead to the deployment of an early product version for an enterprise service ecosystem.

In the medium term, continuous development, deployment, and operational requirements for the capabilities that enable a general-purpose federated compliance intelligence service demand dedicated research efforts within the broader enterprise community. Active research contributions in these areas will therefore spur ecosystem deployment by enabling risk-aware members to migrate from internally developed solutions toward the desired shared enterprise compliance service, thereby improving ROI for the entire community.

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

2. Theoretical Foundations of Federated Compliance Intelligence

The section defines the core concepts—federation, compliance, intelligence, data provenance, trust, and accountability—underlying federated compliance, establishing a foundation for the subsequent infrastructure- and policy-focused research agenda. Compliance intelligence is identified as a federated problem requiring a distributed approach to policy management, cross-domain risk scoring, and joint compliance validation. Finally, a survey of privacy-preserving and secure computation techniques highlights their applicability to federated intelligence, explores current limitations, and positions their future development within the context of data governance and auditability.

Federated compliance intelligence builds on existing frameworks to create a governance layer that oversees collaboration among diverse parties in a zero-trust model. Operationalizing compliance across services and systems from multiple sources entails addressing the research challenges of distributed policy engines, modeling risk across boundaries, and validating compliance within a federated context. Such algorithms broaden the scope of compliance assurances and facilitate collaborative certification. Policy processing across multiple domains supports sharing information assets with confidence, enabling the safe and trusted use of AI-based services built on data from different sources.

Table 2: Layers of Federated Compliance Intelligence Architecture

Layer	Primary Function	Key Components	Compliance Role
Data Layer	Data storage and processing	Data repositories, encrypted storage, provenance logs	Data protection and residency enforcement
Control Layer	Policy governance and authorization	Access controls, policy engines, authorization systems	Enforcement of security and compliance rules
Decision Layer	Risk evaluation and operational decisions	Risk scoring engines, AI analytics, orchestration tools	Automated compliance decision-making
Governance Layer	Cross-domain coordination	Trust frameworks, audit systems, accountability controls	Regulatory alignment and oversight

2.1. Definitions and Core Concepts

Formal definitions of federation, compliance, intelligence, provenance, trust, and accountability support a discussion of compliance as a federated problem incorporating data, control, and decision layers. A federated approach acknowledges the distribution of relevant policies across domains and supports conflict resolution and consensus-building mechanisms.

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

Federation, compliance, intelligence, provenance, trust, and accountability constitute a first-order vocabulary for the present work. These concepts may be further elaborated and detailed through dedicated treatment, but a degree of contextual definition, differentiation, and relationship is necessary to motivate the subsequent development.

Federation enables data-minimization and privacy-preserving architecture patterns at many scales, from cloud-connector brokers exchanging micropayments to service ecosystems sharing telemetric information and regulatory compliance policies. Yet even for informal definitions, federation presents conceptual challenges. The term covers both the organizational distribution of data and the technical properties of decentralized architectures. A federated service ecosystem features independently governed, owner-controlled data stores located across organizational and jurisdictional borders, yet asserts the properties normally associated with a centralized service chain.

The data-management layer addresses the location, protection, management, and availability of data captured, processed, and produced by services within the ecosystem. The control layer manages authorizations and policies related to compliance and risk, whether from the owner, user, or another authority. These policies govern operation of the data layer, allowing enforcement of measures such as read/write restrictions, access scheduling, or user notification. The decision layer operationalizes these policies, incorporating tools for evaluating risk scores and data protection.

Equation 2: Compliance Policy Evaluation Function

The distributed policy engine determines whether an action is permitted or denied.

$$P(a,d) = \begin{cases} 1, & \text{if action } a \text{ satisfies policy in domain } d \\ 0, & \text{otherwise} \end{cases}$$

Where:

- $P(a,d)$= policy validation function
- a= requested action
- d= organizational domain

$$P(a,d) = \begin{cases} 1, & \text{if action } a \text{ satisfies policy in domain } d \\ 0, & \text{otherwise} \end{cases}$$

2.2. Compliance as a Federated Problem

Compliance constitutes an intricate problem structure requiring well-defined semantics to establish compliance policies for disparate regulatory environments and domains. Different parties across distinct locations may find themselves exposed, for instance, to national regulations on data residency or industry-specific regulations such as the Health Insurance Portability and Accountability Act (HIPAA) or the European Payment Services Directive (PSD2). Therefore, a single enterprise cannot bear full responsibility for compliance even when it controls systems that act as federated hubs for on-demand services needing access to the data of other partners. Organizations or individuals running such endpoints inevitably carry the risk that third-party directly or indirectly connecting to their networks may use the underlying information in a way that causes that party to fall out of compliance in their own domains. Residual risk can only be scored and mitigated by performing checks across all partners' territories against the relevant regulations that are in force, using not just their own policy repositories but also those of the other parties involved, something commonly understood as the use of cross-domain risk scoring.

Furthermore, and beyond the obvious need to have the right policies in place at the right time, the design of a compliance task extends to the definition of controls and procedure sets. Procedural models specify sequences of actions to be executed in order to ensure compliance with any given requirement. In a development environment, for instance, such controls can be combined into a continuous compliance approach: checks can be performed for all policies across the full execution flow of the DevOps pipelines. A similar outcome can be defined for SREs in production with a different emphasis, focusing mainly on telemetry for operational-exposed services. Although such models mostly connect the outputs of compliance checks with DevOps and SRE practices, the same holistic orchestration view is valid

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

for incident response—playbooks for handling incidents related to a lack of compliance must address not only actions but also policies.

2.3. Privacy-Preserving and Secure Computation Techniques

Privacy-preserving computation and secure multiparty computation (MPC) techniques enable computations over confidential data without disclosing this data to the parties involved in the computation. Applying these techniques to the compliance domain allows checks to be performed on sensitive information while protecting it from untrusted service operators. Many such techniques have been proposed recently, but they have different applicability, performance, underlying assumptions, and security guarantees, making it essential to identify their respective niches. Several dimensions are relevant for this taxonomy: data ownership (which party holds the data to be processed); the number of parties involved in the computation; the nature of the data sharing (whether sharing of the underlying data, the result of a function computed over this data, or the policies controlling the computation of the function is mandated); the nature of the computation (passive or active); and the security model (the trust level of the parties performing the computation).

Data minimization is a key aspect in the design of services, especially those offered by third parties that have access to data from an ecosystem beyond their own domain. It aims to reduce the amount of data processed by a third-party service to the minimum necessary to achieve the intended objective or a predefined level of acceptable risk. This principle applies not only to the processing of private user data, but also to regulatory requirements. Minimization may be satisfied through keeping data encrypted during processing in such a way that it can be stored in its encrypted form, or through anonymization or pseudonymization techniques. Controls need to guarantee that all incoming data satisfy the conditions for minimization and that these conditions are satisfied at the output for data leaving the ecosystem, including for data used for validation.

Fig 2: Privacy-Preserving Analytics and Secure Multiparty Computation

3. Architectural Paradigms for Scalable Service Ecosystems

Federated architecture patterns capture high-level choices in structuring a compliance ecosystem, ranging from how data ownership is assigned to how the politics of connection negotiation and policy agreement are orchestrated. A service ecosystem is an environment in which heterogeneous services of third parties can interoperate in a distributed way, usually over the Internet. Scaling data-driven service ecosystems so as to enable the venture of each ecosystem participant comes with challenges, such as the related data residency regulations, according to which a data subject's data must be stored and processed in one or more designated countries.

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

Countries want to protect their citizens' privacy by imposing legal obligations with regard to data governance and usage. Often, these regulations require data localization or impose constraints on cross-border flows. Enterprises must comply with legal obligations about deploying and managing services which process data that imply data flows across country borders. Existing solutions abstract away these concerns in order to focus on other aspects of scaling service ecosystems, such as operations and performance, to make service provision more efficient and less expensive. Yet, according to Fylling et al. (2022), legal compliance is also an enabler of scaling service provision within a service ecosystem; it just has not been prioritized.

Table 3: Privacy-Preserving Computation Techniques for Federated Compliance

Technique	Purpose	Advantages	Limitations	Compliance Relevance
Secure Multiparty Computation (MPC)	Compute on distributed confidential data	Strong confidentiality guarantees	High computational overhead	Cross-domain compliance validation
Federated Learning	Train models without centralized data sharing	Data privacy preservation	Non-IID data challenges	Compliance-aware anomaly detection
Anonymization	Irreversible identity removal	Eliminates re-identification risk	Loss of data granularity	GDPR and privacy compliance
Pseudonymization	Replace identifiers with reversible tokens	Maintains analytical usability	Residual re-identification risk	Secure data sharing
Encryption-in-Use	Process encrypted data	Strong data protection	Performance overhead	Sensitive telemetry analysis

3.1. Federated Architecture Patterns

Enterprise service ecosystems offer the promise of rapidly delivering data-rich services while minimizing costs. As the reliance on external services increases, the need to demonstrate compliance with data protection and cybersecurity laws becomes increasingly challenging and costly. Moreover, external services are often monitored by different stakeholders that impose conflicting policies while using different terminology. These violations can incur significant fines for the companies involved, creating a vital need for compliance that scales with the service ecosystem.

Specific architectural patterns used to define services can significantly simplify compliance processes. When services are operated in a hub-and-spoke model, compliance requires satisfying the hub's requirements, while a peer-to-peer model distributes the responsibility of establishing trust among partners. A hybrid model addresses data residency obligations in different jurisdictions.

Equation 3: Federated Anomaly Detection Equation

The runtime monitoring section explains anomaly detection using deviations from expected telemetry values.

$$A(x) = \frac{|\, x - \mu \,|}{\sigma}$$

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

Where:

- $A(x)$= anomaly score
- x= observed telemetry value
- μ= historical mean
- σ= standard deviation

If:

$$A(x) > \tau$$

then an anomaly alert is triggered.

$$A(x) = \frac{| x - \mu |}{\sigma}$$

3.2. Service Mesh and Compliance-Aware Orchestration

The integration of compliance checks into service mesh infrastructure can enable full automation of the compliance check process for enterprise services based on this architectural pattern. Service meshes add a dedicated layer for service-to-service communication over a data path with integrated observability, traffic management, and security capabilities, eliminating the need for implementing these common concerns in each application service. A telemetry data tracker in the service mesh provides insights into the communication between services while relevant security concerns are handled by a sidecar proxy per service instance. The sidecars can extract the relevant parameters defined in the established compliance policy whenever a service calls a different service, and the agent evaluation engine can check whether the service invocation conform to the established policy. When the data passed or received transfer between services contains sensitive data, the data can be anonymized before being transferred.

This orchestration can, thus, be configured to transfer sensitive data only when a compliance violation has not occurred within an established threshold period. Policy drivers extend the service mesh with dedicated components for integrating compliance policies into the service operation. Whenever a new policy is established within the service mesh, the driver dedicated to the sensitive service pops up a message requesting an internal compliance action within the other service being called. Policy-informed routing adds intelligence to the routing decision making in the service mesh, using the telemetry information extracted by the sidecars, so that sensitive information is routed through anonymization services whenever required.

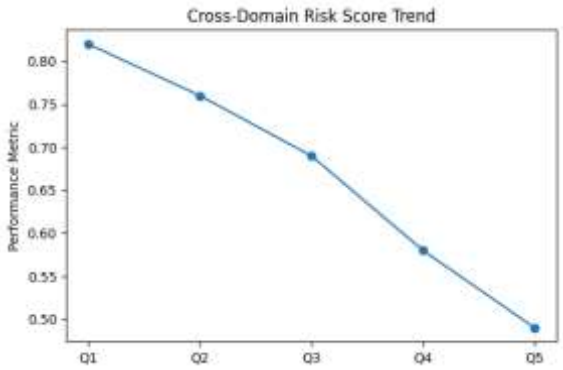

4. Data Governance, Privacy, and Regulatory Alignment

Global or industry-specific laws and regulations impose requirements on established procedures, controls, and technical measures. Most compliance regimes are concerned with specific issues, such as data residency or cross-border data transfer, which are usually addressed in sectoral regulation. Even when different regimes cover similar areas, they are rarely aligned and sometimes even contradict each other. A federated compliance architecture can therefore not ignore these political and regulatory constraints.

Data minimization, data anonymization (irreversibly transforming data so that individuals can longer be associated with the data), and data pseudonymization (replacing the data

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

by tokens that hide the data) are widely recognized as methods to mitigate privacy risks. Each of these methods has advantages and disadvantages, and should therefore be considered carefully for each case. Mitigation or removal of several risks will also enhance acceptance of data usage for other purposes. Criteria should be established to guide which method to adopt for which data elements, and an explicit definition of the risk threshold should be included, since it will have a strong impact on the choice of operation. Data retention policies should be designed, specifying how long data elements should be kept at minimum and maximum as well as the conditions under which they may be kept longer.

Table 4: Federated Architecture Patterns and Compliance Implications

Architecture Pattern	Description	Advantages	Compliance Challenges	Suitable Scenario
Hub-and-Spoke	Central hub coordinates services	Centralized governance	Single-point compliance dependency	Enterprise cloud ecosystems
Peer-to-Peer	Direct collaboration among services	Distributed trust and resilience	Complex trust negotiation	Cross-organizational ecosystems
Hybrid Federated Model	Combination of centralized and decentralized operations	Flexible jurisdiction handling	Policy synchronization complexity	Multi-country enterprises
Service Mesh Architecture	Sidecar-based service communication	Integrated telemetry and security	Operational complexity	Microservices environments

4.1. Global and Industry-Specific Compliance Requirements

Compliance requirements may stem from diverse regulations at global, national, and sectoral levels. For instance, the EU's GDPR mandates data Residency and transfer rules for personal information, while market-entry regulations might impose specific storage requirements for foreign digital service providers. Industry regulation further complicates matters, introducing requirements about providing specific functionalities or obtaining prior approvals for certain decisions. Despite these differences, many regimes share principles such as minimizing data collection or processing. Complexity arises at system operation or service usage, where several requirements must be met simultaneously, and at cross-border services relying on another domain, where requirements might diverge. Given the interconnectedness of social, economic, and legal systems, compliance outcomes must often be jointly considered with partners across domains.

These issues must be addressed through the entire ecosystem lifecycle: when operationalizing a system, designing it, or using it. Therefore, regulatory requirements must be mapped to architectural primitives capable of satisfying them, and consideration must be given to methodologies for enforcement and functioning-altering aspects.

Equation 4: Privacy-Preserving Data Minimization Function

The paper emphasizes minimizing data exposure in federated systems.

$$D_{min} = \arg\min_{D'}(Risk(D'))$$

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

Subject to:

$$Utility(D') \geq U_{threshold}$$

Where:

- D'= reduced dataset
- $Risk(D')$= privacy exposure risk
- $Utility(D')$= operational usefulness of data

$$D_{min} = \arg\min_{D'}(Risk(D'))$$

4.2. Data Minimization, Anonymization, and Pseudonymization Strategies

Data minimization reduces risks and costs associated with data storage and processing. It limits retention to the minimum necessary for business operations, in line with regulatory guidelines. Removal of data no longer in use serves business, security, and compliance interests. Organizations impact their risk profile through policies governing the storage and sharing of data, the time frame for anonymizing or destroying data after completion of its intended use, and controls for ensuring appropriate data handling by third parties. Anonymization—defined as irreversible data transformation that prevents identification and reidentification—is an effective way to eliminate risk while maintaining the ability to analyze themes or trends. Such data can be retained for extended periods at low cost and shared with a high degree of confidence.

Even though anonymized data cannot be traced back to its originating entity, decentralized storage and sharing significantly reduce reidentification risk. The adherence to best practices, however, still requires organizations to enforce strict disclosure controls. Pseudonymization (for reversible de-identification) and generalization (controlling the granularity of information) are other accepted privacy-preserving methods. These latter approaches reduce rather than eliminate reidentification risk and thus necessitate the enforcement of stricter disclosure controls. Anonymization, however, is not always achievable. In these cases, proper use of resources and mitigation of risk is paramount. An organization may be compelled to use a third-party service and purchase data; in which case, it is required to mask identifiers for personal data that are not under its direct control while sharing with the service provider.

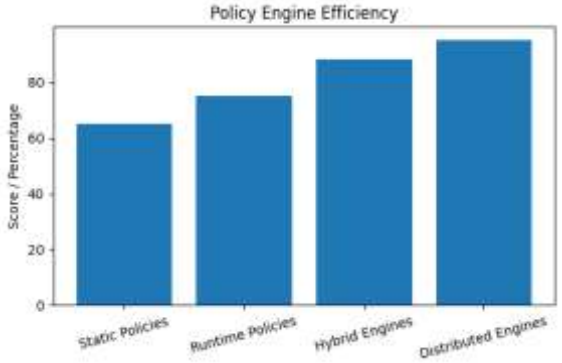

4.3. Auditability, accountability, and Transparency Mechanisms

Audit trails, immutable logs, and decision explainability aid post-factum assessments, provide stakeholders with security guarantees, and establish trust in automated systems. Auditability examines whether a system generates sufficient evidence for a trustworthy evaluation of actions and decisions, and is often assessed using compliance checks against pre-defined policies and contracts. Auditing emphasizes an external perspective. In enterprise service ecosystems, the high degree of automation and autonomicity necessitate continuous compliance assessment and facilitate the definition of near-real-time auditability metrics. Retention policies define the conditions and duration for maintaining audit data, including the principles of data minimization, anonymization, pseudonymization, and data absenteeism.

Accountability concentrates on the determination of responsibility and liability. It is particularly complicated for autonomous systems, where monitoring and supervision by

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

humans is not feasible. A trust and accountability framework maps the interaction between each involved party and specifies the conditions under which trust is appropriately established. Transparency represents the ability of its internal state and behavior to be comprehensible. It is crucial in the context of algorithms, where neural networks are often regarded as black boxes that can make decisions relying on hidden logic. It can also be achieved by providing posthoc justifications for decisions. Key aspects of auditability, accountability, and transparency applied in enterprise service ecosystems are the definition of audit trails, the use of tamper-proof logs, and the explainability of decisions made by algorithms.

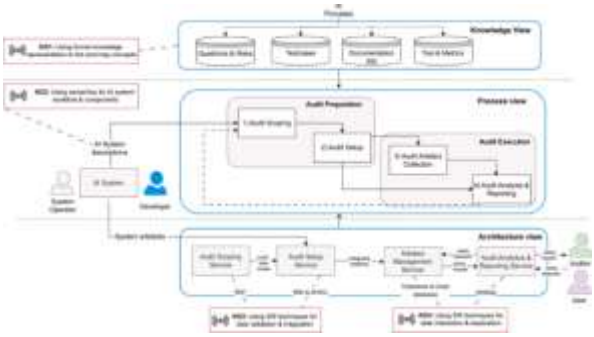

Fig 3: Auditability, accountability, and Transparency Mechanisms

5. Federated Compliance Intelligence: Methods and Algorithms

Compliance validation in federated ecosystems requires diverse techniques to address a wide range of federated compliance-checking problems. Distributed policy engines validate compliance against domain-specific policies, while cross-domain risk scoring assesses the safety of sharing or combining sensitive information across different contexts and organizational boundaries. Model-driven methods enable compliance validation for complex regulatory frameworks spanning multiple domains.

Policies defined across multiple domains influence the nature of relationships between ecosystem partners and establish conditions for safe data sharing and collaboration. These policies represent the governance layer for federated ecosystems and ensure adherence to regulations. Distributed policy engines are designed for compliance validation within such contexts, encompassing policy definition and deployment, deployment across domains, automated conflict resolution, and execution. The specification of input and output models and policy languages, along with the semantics of policy evaluation, is crucial. Governance influences the design of distributed policy engines by defining the trust relationships among policy owners and supporting voting mechanisms. As policies become more general and express higher-dimensional connections, performance and security considerations assume greater importance.

Risk scoring assesses compliance and safety when sharing or combining sensitive information across multiple domains. It operates through the aggregation, normalization, and explanation of risk scores originating from several data providers and/or in different contexts. A method for cross-domain risk scoring is needed to determine whether it is safe to share or combine sensitive information.

Table 5: Distributed Policy Engine Components

Component	Function	Example Activities
Policy Repository	Stores compliance policies	GDPR, HIPAA, PSD2 rules
Runtime Evaluation Engine	Evaluates requests against policies	Service invocation validation
Sidecar Proxies	Extract telemetry and policy parameters	Traffic inspection and anonymization
Conflict Resolution Module	Resolves cross-domain policy conflicts	Consensus generation

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

Component	Function	Example Activities
Synchronization Manager	Maintains policy consistency	Policy updates across regions
Audit Logging System	Tracks policy evaluation outcomes	Compliance evidence generation

5.1. Distributed Policy Engines

Distributed policy engines answer the question "Do the policies of a domain or organization allow a particular action to be performed?" at any layer of the architecture, such as data ingress, data egress, and service invocation. They allow policies to be statically stored on management nodes (controllers, management plane) or conditionally evaluated in an ad-hoc manner during runtime (service-mesh sidecars, data-saving during processing) through syntactically and semantically compatible security policy languages. Their inputs are domain-dependent and cover all types, including data sources, actions such as transferring and using data, vulnerable components located in the service mesh, paths of the data flow, and usage or transfer requests. The outputs typically indicate whether permission is granted or denied.

Policies stored on management nodes (e.g., ingress/egress routers, service mesh controllers) or run-time engines (e.g., sidecars) must be periodically synchronized. Conflicts among region- or domain-specific policies should be defined and resolved as part of a policy lifecycle management process. A dedicated policy engine that participates in establishing semantic consensus among local policy evaluation engines and guarantees that those engines' outputs/Cd are consistent could be useful. Distribution of policy engines enhances assurance through redundancy; their deployment in high-risk components or regions is therefore preferred.

Equation 5: Federated Trust Score Equation

The trust framework discussed in governance and accountability sections can be represented as:

$$T_e = \alpha C + \beta S + \gamma A$$

Where:

- T_e= enterprise trust score
- C= compliance score
- S= security posture score
- A= auditability score
- α, β, γ= weighting coefficients

with:

$$\alpha + \beta + \gamma = 1$$

$$T_e = \alpha C + \beta S + \gamma A$$

5.2. Cross-Domain Risk Scoring

Risk scoring identifies the vulnerability of an entity to a specific form of risk, emphasizing its preparedness to withstand threats and covert or overt attempts to exploit an inherent weakness. Each evaluation action is rooted in a particular source of information, addressing a well-defined aspect of compliance, and accounting for broad environmental features. The results can be supervised by domain controllers or policy makers and then aggregated to provide a complete feature set for subsequent decisions or actions. Such risk numerical outcomes are most useful when prepared by services of trusted and known domains.

Risk scoring across different domains requires careful aggregation of individual scores to obtain an input for a process driven by another domain. Normalizing each score within its domain and risk category is an important prerequisite for achieving a meaningful overall risk assessment. Labeling the scores according to the domain, potential path or actor, and risk category facilitates the creation of an explainable final assessment.

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

Normalizing the scores between 0 and 1 before aggregation clarifies the ultimate assessment's meaning. Specific paths among different domains can also be examined. For instance, if an exposed application is eventually operating within domain A, receives some information or help from service B, and can patrol its perimeter by querying domain C, the official crowd-sourcing risk-score output can be shaped as the maximum value of the three services with respect to the 'exploitation' risk, the minimum value with respect to the 'resilience' risk, and the average with respect to the 'oversight' risk.

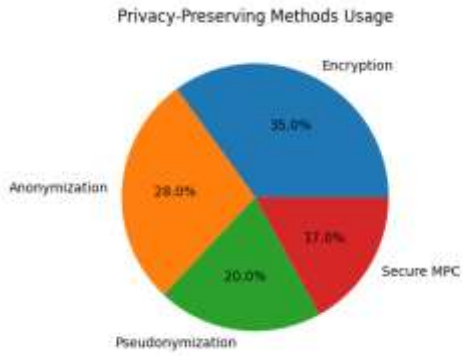

5.3. Model-Driven Compliance Validation

Model-driven compliance validation leverages formal models to ascertain compliance across evolving service ecosystems. Modelling languages, such as SysML and BPMN, can depict property requirements and system implementations, enabling consistency checks and simulations against internal and external regulations. Furthermore, interface definitions facilitate the creation of model-driven validation pipelines.

Policy inconsistencies and compliance violations remain a challenge during the development of enterprise service ecosystems. Model-driven approaches for validating properties across complex software systems are well-established. Property requirements can be succinctly expressed in selected modelling languages, such as SysML and BPMN. These languages can also describe the implementation of enterprise service ecosystems, supporting simulation to verify conformity to the required properties. Even when properties are not preserved due to implementation changes (e.g., adding a new bidirectional channel between two domains), violation information can be beneficial for stakeholders.

Requirements can be expressed as SysML or BPMN invariants. BPMN interface definitions enable simulation to compare the implementation with the properties and to build compliance-validation pipelines. Although property violation cannot be entirely eliminated, a model-driven pipeline for SysML-to-BPMN transformation, simulation, and compliance validation can help owners of an enterprise service ecosystem gain confidence in their operations. The pipeline can be defined as an abstraction hierarchy, with the lower layers designed to support its economy and coverage without introducing excessive complexity.

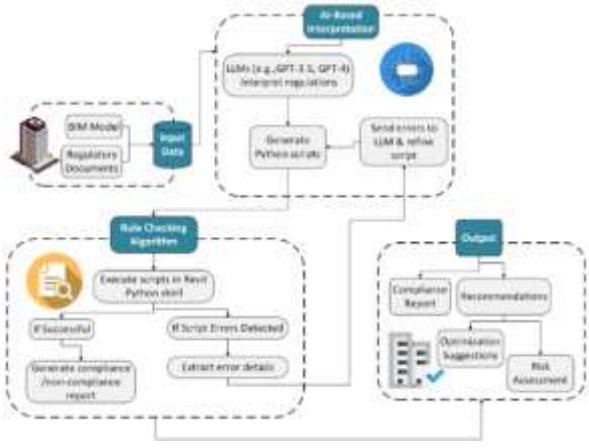

Fig 4: Model-Driven Compliance Validation

6. Operationalization in Enterprise Service Ecosystems

The practical integration of compliance and governance practices in enterprise service ecosystems is examined. The

focus is placed on DevOps and Site Reliability Engineering pipelines, where security concerns are embedded into the evolution process. Guardrails are defined to prevent critical violations, while violations of non-critical policies are gradually remediated. Telemetry from the runtime environment can be used to monitor compliance, detect anomalies, and support incident response and forensics. Three risk-related aspects are specified: trusted application telemetry, clear procedures for cross-domain incident response, and playbooks to preserve forensic evidence.

Compliance concerns are becoming increasingly important during the implementation and operation of enterprise service ecosystems. Service-failure incidents can have long-lasting reputational impact, making cyber incidents more frequent, damaging, and costly. Security best practices are not globally applied, supported, or enforced; tangible telemetry on the operational state of services is rarely analyzed and structured. To address these issues, compliance checks need to be integrated into DevOps and SRE processes through appropriate guardrails, ensuring that non-compliance is simply the result of an operational backlog. In addition, security incidents that cross-domain and trust boundaries must have well-defined procedures to allow for a proper incident response.

Table 6: Cross-Domain Risk Scoring Framework

Risk Dimension	Description	Example Metric	Aggregation Method
Exploitation Risk	Vulnerability exposure level	Threat probability	Maximum score
Resilience Risk	Ability to withstand attacks	Recovery readiness	Minimum score
Oversight Risk	Governance effectiveness	Audit coverage	Average score
Trust Risk	Reliability of ecosystem partner	Trustworthiness index	Weighted aggregation
Privacy Risk	Potential privacy violations	Data sensitivity score	Normalized scoring

6.1. Continuous Compliance in DevOps and SRE

The goal of continuous compliance is to prevent violations in production by embedding policy checks in DevOps and Site Reliability Engineering (SRE) processes. A continuous compliance pipeline augments existing CI/CD pipelines with checks against appropriate policies, ensuring that no policy violations are introduced by code committed at any point in time. Policy checks may include static code analysis (linting), dynamic security scans, tests against production operation policies, and others. These policy checks can also monitor production as part of SRE practices, capturing events in the operations database, validating them against operational policies, and raising incidents for relevant stakeholders if any deviation occurs.

Continuous compliance implementations depend on specific integrations with CI/CD, telemetry, and incident management systems. All integration points must be captured in a compliance pipeline for comprehensive auditing, detection velocity assessment, and analysis of operational bliss violations over time. To maintain operational effectiveness, changes that violate controls must also offer a rollback strategy to the last compliant state. Continuous compliance speed is equally critical: controls need to be expressed as code so compliance velocity can be measured and reduced over time.

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

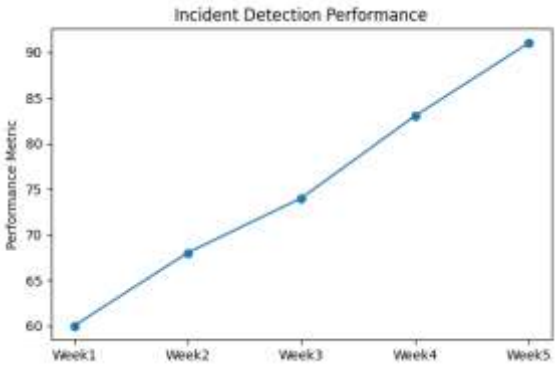

6.2. Runtime Monitoring and Anomaly Detection

Telemetry collection and analysis are prerequisites for incident detection and response. Anomaly detection models trained on historical data can trigger alerts when incoming telemetry values deviate significantly from predicted ranges. Recent research demonstrates that such models can be effectively and efficiently implemented in a federated setting, where training data are distributed across different organizations, preserving their privacy and confidentiality. Nevertheless, the definition of an anomaly for any single organization depends on the external context in which it operates. For instance, the fact that the maximum number of CPU cores in use at any moment of time during the last days was 80 cannot by itself be considered an anomaly; any future value equal to 80 should not trigger any alert. Malicious or anomalous behaviors are typically characterized by small probabilities of occurrence given the training data. As a consequence, the development of appropriate ground truth data for training and testing anomaly detection models is challenging. Moreover, organizations have different privacy and legal agreements, which may forbid them to voluntarily share information.

These considerations highlight the usefulness of federated learning for telemetry data. Such federated solutions can optimize the use of local data to improve models detecting malfunctions or anomalies in a partner's infrastructure. The trained model can then be evaluated against a specific organization's robustness to sensitive information, so that the inspection of incoming telemetry is comparable to when the events occurred in the training dataset. This process can also be extended to include other detected anomalies, such as increased operating temperatures or abnormal power consumption, within the model training, so that it approximates better the domain being used.

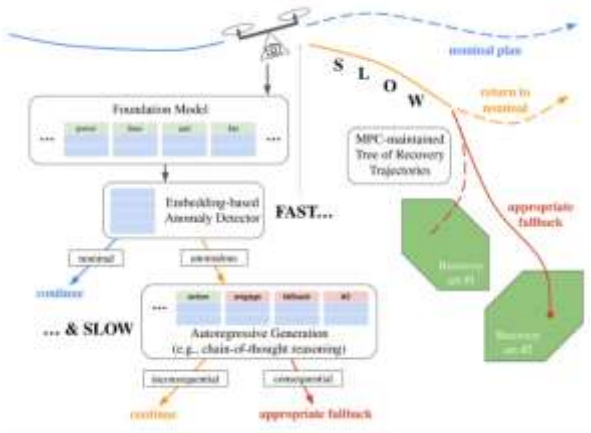

Fig 5: Runtime Monitoring and Anomaly Detection

6.3. Incident Response and Forensics in Federated Contexts

Thorough preparations and timely response decrease the negative impact of security incidents. Initiating effective incident response procedures require special training for a dedicated team, defining the relevant playbooks, and establishing the required people, processes, and technology. Some elements of incident response can be automated: if available, dedicated platforms can collect and correlate telemetry, and process rules can recommend simple responses, such as blacklisting a malicious IP. Even in those cases where a human must decide on the right course of action, having accessible and well-documented procedures helps in striving through the incident. Investigation is part of every incident response; ideally, the designated team verifies compliance and forensics rules beforehand.

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

In enterprise service ecosystems, the collaborative nature of the domains involved and the partners involved necessitate the definition of dedicated playbooks for the discovery, reporting, investigation, and resolution of incidents that span multiple domains. These processes must keep compliance and forensics services in their minds and define how to capture and preserve the required incident evidence without violating applicable policies, respecting data privacy, and enforcing data minimization. The earlier incident data are collected, the smaller the risk of the evidence being tampered with—or even lost altogether. Organizations should also ensure that the relevant partners are alerted as soon as possible, so they may start gathering evidence that is under their control.

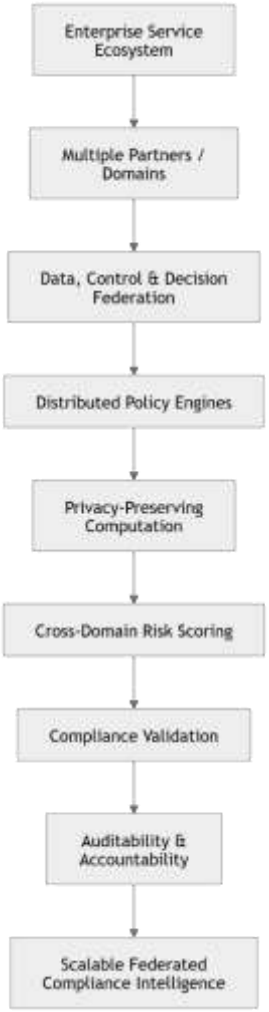

Federated Compliance Intelligence Architecture

7. Governance, Risk, and Compliance (GRC) Architecture

A GRC architecture defines roles, responsibilities, and accountability chains for governance, risk, and compliance. Proper separation of duties and application of access-control

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

mechanisms prevent collusion and support detection of anomalous activities.

Governance, risk management, and compliance functions define who is responsible for what, and how different roles formally relate to each other. The GRC architecture specifies for every function that relies on information-sharing among multiple parties: Who enforces which policy? Whose data should be collected or contributed in what circumstances? Who can access what information for what purpose? Who is in charge of responding to which type of incident?

Failure to answer these questions properly can allow one person to bypass, undermine, or ignore multiple controls. Strong separation of duties prevents collusion: ensuring that, whenever possible, two or more distinct individuals are required to complete a critical task. When separation of duties is impractical, role-based access controls provide an alternative mechanism: denying certain roles access to certain functions in a way that prevents any single role from prying into others' data or abusing its power.

7.1. Roles, Responsibilities, and Accountability Chains

The governance architecture identifies roles, responsibilities, and accountability chains within enterprise service ecosystems. A variety of participants are involved—each with specific duties—contributing to the maintenance of the operating environment. The inclusion of internal teams requires careful separation of duties to mitigate collusion risk, while external stakeholders depend on external auditor evaluation, certification, and testing to foster trust.

As enterprise service ecosystems adopt a federated model, efforts to compress involved parties' administrative operations slacken, and the role, aGraph, context, and degree of responsibility for specific policies and processes become paramount. Figure 7 provides a simplified overview, positioning the executive at the top of the operational hierarchy, underlying and delegating day-to-day operations through junior roles within the governance and strategy functions. The role is responsible for defining the actors' expectations, aligning with non-federated environments; thus, creating an institutional role to serve as a strategic business unit—known as comgovern—dedicated to governance activities provides a feasible temporal scheme for fulfilling ecosystem needs.

As the number of roles increases, the allocation of specific duties to dedicated bodies becomes a feasible strategy. In the absence of roles specializing in a specific area, the overall distribution of duties requires packaging policies and responsibilities into concise, understandable clusters, allowing for the definition of policies rather than isolated actions. The indeterminate transit, transit-role, and other security policies, implemented as a single transport policy, serve as an example.

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

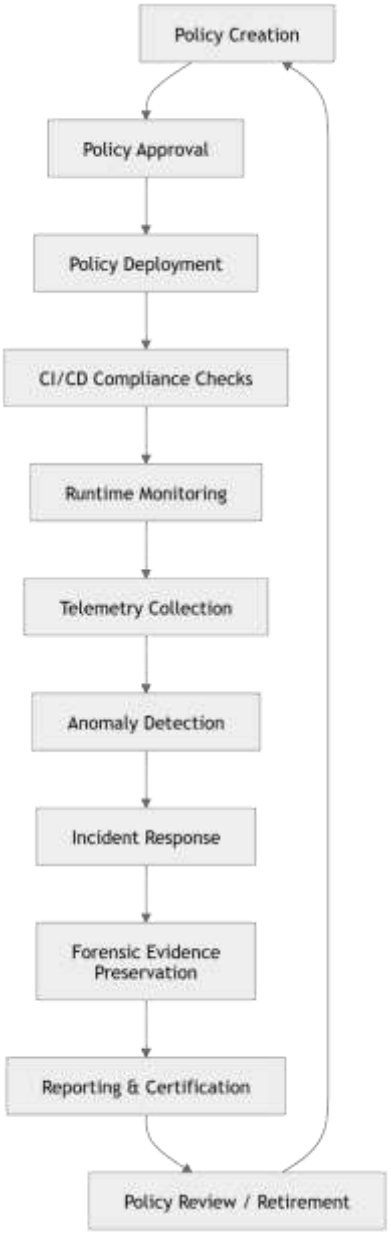

Operational Compliance Lifecycle

7.2. Policy Lifecycle Management

Policies guide operational actions and transactions. The lifecycle of policies includes their creation, amendment, versioning, and retirement. Authors initiate the process, which requires approval from designated owners. Once in effect, the policies are deployed to operational systems. Metrics ascertaining the effectiveness of existing policies can lead to a mandate for amendment. Furthermore, when external conditions enforce a new policy requirement, the creation of a new policy is triggered. Policies may also be marked for retirement, in which case existing assignments need to be validated and undertaken by relevant parties. The final process involved in the retirement of policies includes the deployment of alerts on upcoming retirements.

Policy management encompasses the management of policy creation, approval, deployment, revision, and retirement. A formal policy management specification specifies the support for versioning and traceability. Policy layer management within governance risk and compliance enforces separation of duties. The corresponding access control model mandates access to the policy creation function to a single identified creator. The access control mandates that the approving function is executed by a set of approving authorities and that policy deployments, revisions and retirements are performed by control self-assessment owners.

8. Conclusion

Federated Compliance Intelligence for Scalable Enterprise Service Ecosystems: All components of a scalable enterprise service ecosystem are in place, but technical and organizational progress has stalled—data protection, cybersecurity, and fraud-prevention compliance controls remain largely independent and highly localized. Both practical implementation and validation of the proposed principles, patterns, and methods are essential; real-world pilots should advance compliance-in-the-large and be supported by further development of standards for privacy-preserving and secure computation in federated contexts.

Recent research has focused on a specific aspect of this challenge: federated compliance intelligence, the provision of automated support for compliance-related decisions across contextual boundaries.

A federated lens on compliance enables consideration of the many objectives beyond privacy that enterprises must address both individually and collectively. Compliance directives of all kinds often coalesce into requirements for risk scoring, preparatory documentation, incident response, and auditability. Automating these processes at scale requires constant attention, but compliance still cannot be regarded as fundamentally different from security—neglect invites disaster, but systematic investment alone does not ensure immunity.

References

1. Bonawitz, K., Kairouz, P., McMahan, H. B., Ramage, D., & others. (2022). Federated learning and privacy. Communications of the ACM, 65(4), 90–97.
2. Mothukuri, V., Meimandi Parizi, R., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. Future Generation Computer Systems, 115, 619–640.
3. Kolla, T. (2024). Graph Neural Networks for HCC Risk Adjustment and Interoperability. International Journal of Science, Research and Technology, 7(6), 13244-13255.
4. Yin, X., Zhu, Y., & Hu, J. (2021). A comprehensive survey of privacy-preserving federated learning: A taxonomy, review, and future directions. ACM Computing Surveys, 54(6), 1–36.
5. Zhu, H., Xu, J., Liu, S., & Jin, Y. (2021). Federated learning on non-IID data: A survey. Neurocomputing, 465, 371–390.
6. Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., & Gao, Y. (2021). A survey on federated learning. Knowledge-Based Systems, 216, 106775.
7. Inala, R. Designing Scalable Technology Architectures for Customer Data in Group Insurance and Investment Platforms.
8. Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., & Poor, H. V. (2021). Federated learning for Internet of Things: A comprehensive survey. IEEE Communications Surveys & Tutorials, 23(3), 1622–1658.
9. Nguyen, D. C., Pham, Q.-V., Pathirana, P. N., Ding, M., Seneviratne, A., Lin, Z., Dobre, O. A., & Hwang, W.-J. (2022). Federated learning for smart healthcare: A survey. ACM Computing Surveys, 55(3), 1–37.
10. Zhang, K., Song, X., Zhang, C., & Yu, S. (2022). Challenges and future directions of secure federated learning: A survey. Frontiers of Computer Science, 16, Article 165817.
11. El Ouadrhiri, A., & Abdelhadi, A. (2022). Differential privacy for deep and federated learning: A survey. IEEE Access, 10, 22359–22380.
12. Davuluri, P. S. L. (2023). AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems (December 15, 2023).
13. Qu, Y., Uddin, M. P., Gan, C., Xiang, Y., Gao, L., & Yearwood, J. (2022). Blockchain-enabled federated learning: A survey. ACM Computing Surveys, 55(3), 1–35.
14. Huang, J., Kong, L., Chen, G., Xiang, Q., & others. (2022). Blockchain-based federated learning: A systematic survey. IEEE Network, 36(6), 198–205.
15. Ma, X., Zhu, J., Lin, Z., Chen, S., & Qin, Y. (2022). A state-of-the-art survey on solving non-IID data in federated learning. Future Generation Computer Systems, 135, 244–258.
16. Inala, R. (2023). AI-powered investment decision support systems: Building smart data products with embedded governance controls. Journal for ReAttach Therapy and Developmental Diversities, 6(10), 2251-2266.
17. Hosseinzadeh, M., Rahmani, A. M., & others. (2022). Federated learning-based IoT: A systematic literature

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 01
RECEIVED: JANUARY 27 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 13

review. International Journal of Communication Systems, 35(11), e5185.

18. Nassef, O., Sun, W., Purmehdi, H., Tatipamula, M., & Mahmoodi, T. (2022). A survey: Distributed machine learning for 5G and beyond. Computer Networks, 207, 108820.
19. Soykan, E. U., Karaçay, L., Karakoç, F., & Tomur, E. (2022). A survey and guideline on privacy enhancing technologies for collaborative machine learning. IEEE Access, 10, 95719–95743.
20. Kolla, S. K., & Mangalampalli, B. M. (2024). Edge-Based Deep Learning Systems for Point-of-Care Diagnostic Intelligence. Journal of Neonatal Surgery, 13(1), 2387-2399.
21. Mökander, J., & Floridi, L. (2022). From algorithmic accountability to digital governance. Nature Machine Intelligence, 4, 508–509.
22. Rafi, T. H., Noor, F. A., Hussain, T., & Chae, D.-K. (2023). Fairness and privacy preserving in federated learning: A survey. Information Fusion, 95, 1–17.
23. Coelho, K. K., Nogueira, M., Vieira, A. B., Silva, E. F., & Nacif, J. A. M. (2023). A survey on federated learning for security and privacy in healthcare applications. Computer Communications, 207, 113–127.
24. Kolla, S. K., & Reddy, V. A. R. (2024). Evaluating Cloud-Native vs. Hybrid Architectures for Health Benefit Administration Systems. International Journal of Medical Toxicology and Legal Medicine, 27(5), 1042-1053.
25. Qammar, A., Karim, A., Ning, H., & Ding, J. (2023). Securing federated learning with blockchain: A systematic literature review. Artificial Intelligence Review, 56, 3951–3985.
26. Martínez Beltrán, E. T., Quiles Pérez, M., Sánchez Sánchez, P. M., López Bernal, S., Bovet, G., Gil Pérez, M., Martínez Pérez, G., & Huertas Celdrán, A. (2023). Decentralized federated learning: Fundamentals, state of the art, frameworks, trends, and challenges. IEEE Communications Surveys & Tutorials, 25(4), 2983–3013.
27. Tan, A. Z., Yu, H., Cui, L., & Yang, Q. (2023). Towards personalized federated learning. IEEE Transactions on Neural Networks and Learning Systems, 34(12), 9587–9603.
28. Ye, M., Fang, X., Du, B., Yuen, P. C., & Tao, D. (2023). Heterogeneous federated learning: State-of-the-art and research challenges. ACM Computing Surveys, 56(3), Article 79, 1–44.
29. Sirohi, D., Kumar, N., Rana, P. S., Tanwar, S., Iqbal, R., & Hijjii, M. (2023). Federated learning for 6G-enabled secure communication systems: A comprehensive survey. Artificial Intelligence Review, 56, 12651–12743.
30. Gottimukkala, V. R. R. (2024). Federated Learning Approaches for Fraud Detection in International Payment Systems. https://www. jisem-journal. com/download/118_JISEM. pdf.
31. Tang, Y., Zhang, Y., Niu, T., Li, Z., Zhang, Z., & others. (2024). A survey on blockchain-based federated learning: Categorization, application and analysis. Computers, Materials & Continua, 79(3), 1–30.
32. Qi, P., Chiaro, D., Guzzo, A., Ianni, M., Fortino, G., & Piccialli, F. (2024). Model aggregation techniques in federated learning: A comprehensive survey. Future Generation Computer Systems, 150, 272–293.
33. Huang, W., Ye, M., Shi, Z., Wan, G., Li, H., Du, B., & Yang, Q. (2024). Federated learning for generalization, robustness, fairness: A survey and benchmark. IEEE Transactions on Pattern Analysis and Machine Intelligence, 46(12), 9387–9406.
34. Zhang, Y., Zeng, D., Luo, J., Fu, X., Chen, G., Xu, Z., & King, I. (2024). A survey of trustworthy federated learning: Issues, solutions, and challenges. ACM Transactions on Intelligent Systems and Technology, 15(6), Article 112, 1–47.
35. Inala, R., & Somu, B. (2024). Agentic ai in retail banking: Redefining customer service and financial decision-making. Journal of Artificial Intelligence and Big Data Disciplines, 1(1), 1-19.
36. Saha, S., Hota, A., Chattopadhyay, A. K., Nag, A., & Nandi, S. (2024). A multifaceted survey on privacy preservation of federated learning: Progress, challenges, and opportunities. Artificial Intelligence Review, 57, Article 184.

37. Han, Q., Lu, S., Wang, W., Qu, H., Li, J., & Gao, Y. (2024). Privacy preserving and secure robust federated learning: A survey. Concurrency and Computation: Practice and Experience, 36(13), e8084.
38. Attard-Frost, B., Brandusescu, A., & Lyons, K. (2024). The governance of artificial intelligence in Canada: Findings and opportunities from a review of 84 AI governance initiatives. Government Information Quarterly, 41, 101929.