Programmable Compliance

Compliance Architecture for Regulated Tokenised Assets

June 2026

Disclaimer

This report and its contents are made available on an “as-is” basis without warranties of any kind. The content in this report does not constitute regulatory, financial, legal or any other professional advice and should not be acted on as such.

None of its authors and contributors shall be liable for any damage or loss of any kind howsoever caused as a result of the use of the information contained or referenced in this report.

Contents

Introduction

Interest in the use of distributed ledger technologies (DLT) for regulated financial activities has grown significantly in recent years. Financial institutions, market infrastructures, and public sector initiatives are increasingly exploring the issuance, transfer, and settlement of tokenised assets, such as fixed income instruments, deposits, and fund units. These developments are driven by the potential for improved efficiency, programmability, and composability in financial markets. That said, a fundamental premise remains: regulated financial activity conducted using distributed ledger technologies must meet the same legal, regulatory and supervisory expectations as in traditional systems.

Translating compliance objectives and controls to programmable environments is, however, non-trivial. Tokenised asset systems, which typically run on DLT, execute transactions automatically and consistently across all participants. This requires institutions to encode the rules governing participation and permissible activity into logic that operates across multiple platforms and participants. This makes it harder to preserve a clear separation between policy intent, operational processes, and technical implementation.

Current implementations, however remain heterogeneous and tightly coupled to specific platforms, asset types, or institutional arrangements. This limits consistent application of compliance checks across use cases and creates challenges for interoperability, particularly where assets move across systems or jurisdictions. It also leads to duplicative implementation, as institutions may need to rebuild compliance processes for each platform. This makes it difficult to assess whether counterparties have performed the necessary compliance checks or to rely on those outcomes with confidence.

Embedding compliance controls directly into asset tokens[1], so that transfer conditions are enforced at the point of execution rather than through separate institutional processes, brings compliance closer to the point of transfer initiation, but creates limitations that become significant at scale. Rules fixed at deployment can only be updated by modifying the contract itself, which becomes increasingly

[1] The use of compliance controls embedded at the token level—such as whitelisting of wallet addresses or closed-loop transfer restrictions—is among the approaches being considered in emerging regulatory frameworks for digital assets. The Hong Kong Monetary Authority's Anti-Money Laundering and Counter-Financing of Terrorism (AML/CFT) framework for licensed stablecoin issuers under the Stablecoins Ordinance identifies this as one option for satisfying applicable compliance obligations alongside other permissible approaches (Hong Kong Monetary Authority, 2025). The European Union (EU)'s framework for asset-referenced and e-money tokens under Markets in Crypto-Assets Regulation (MiCA), together with the application of Travel Rule requirements to crypto-asset transfers under the Transfer of Funds Regulation, similarly create conditions in which technical controls at the token or transfer level represent one practical means of meeting applicable requirements (European Parliament and Council of the European Union, 2023).

burdensome as regulatory requirements evolve. This approach also does not readily address how the obligations of regulated institutions that handle assets issued by others, including settlement agents, custodians, and operators, can be fulfilled through their own controls. Nor does it provide a mechanism for applying different compliance conditions depending on the counterparties, corridors, or jurisdictions involved in a given transaction[3].

This paper proposes an architectural model that addresses these limitations by separating policy definition, identity management, compliance evaluation, and transaction execution into distinct but coordinated components. The separation of policy from execution means compliance requirements can be updated through governed policy changes without redesigning deployed assets. The independence of the execution layer from the base asset allows regulated institutions to enforce their applicable obligations on assets they do not issue. The coordination of compliance evaluation through a common orchestration layer enables checks to be tailored to the specific counterparties, jurisdictions, and use case of each transaction rather than applied uniformly.

The architecture standardises compliance information and uses a Compliance Attestation to record the results of compliance checks, creating a common format that allows institutions to verify counterparty compliance without custom integrations. This shared approach also supports the development of consistent compliance standards across platforms and jurisdictions.

The Global Layer One (GL1) initiative is seeking to build on this foundation by exploring how such architectures can inform the development of shared standards and support their application across financial use cases. This reflects a broader effort to enable programmable financial infrastructure to operate in a manner that is consistent with regulatory expectations while remaining adaptable to technological and market evolution.

[3] For example, the variation in how AML/CFT standards for virtual assets are applied across jurisdictions is documented in Financial Action Task Force (FATF)'s periodic implementation reviews. As of 2024, 75% of assessed jurisdictions were only partially or not compliant with Recommendation 15 on virtual assets, with only one jurisdiction achieving full compliance. Of 65 jurisdictions that had enacted Travel Rule legislation, only 17 had taken supervisory or enforcement action focused on Travel Rule compliance (FATF, 2024).

Existing Programmable Compliance Efforts

Efforts to operationalise programmable compliance have accelerated alongside the growth of tokenised financial assets and distributed ledger infrastructures. Several broad approaches have emerged across industry, standard-setting bodies, and public sector experiments, addressing different layers of the compliance stack: execution-level controls, attestations of checks performed outside the distributed ledger environment, verifiable identity, and composability across actors and jurisdictions.

This section reviews four broad approaches: (1) embedding enforceable controls without exposing sensitive data, (2) attaching verifiable representations of off-chain compliance processes, (3) establishing verifiable identity layers, and (4) enabling composable compliance across distributed actors.

Approach 1: Enforcing Controls While Minimising On-Chain Data Exposure

A core design challenge in programmable compliance is how to enforce regulatory constraints such as transfer restrictions, eligibility rules, and administrative interventions without placing sensitive personal or institutional data on distributed ledgers.

Several initiatives address this by executing compliance logic at the transaction or token level, while keeping underlying compliance data outside the ledger and under institutional control. Work by Kinexys by J.P. Morgan and MIT Digital Currency Initiative explore design patterns for embedding policy constraints into token behaviour while minimising data exposure[4]. The Global Blockchain Business Council's Capital Markets Risk Mitigation Framework similarly outlines principles for separating compliance validation processes from on-chain execution[5].

The shared principle across these initiatives is that compliance outcomes, such as whether a transfer is permitted, can be enforced within the transaction layer, while the inputs to those decisions remain governed within regulated domains. This

[4] The paper examines payment token design patterns centred on safety, integrity, interoperability, and usability (Toh et al., 2025). Its relevance here lies in showing how token-level policy constraints can be combined with architectures that minimise exposure of sensitive compliance data.
[5] The framework focuses on non-financial risks in blockchain infrastructure and highlights architectural approaches for separating validation and control processes from core execution layers (Global Blockchain Business Council & Oliver Wyman, 2026).

supports privacy, aligns with data protection requirements, and preserves institutional accountability over sensitive information.

Approach 2: Verifiable Representations of Compliance Processes Conducted Outside of Tokenised Asset Environments

A second pattern builds on this separation by representing compliance outcomes in a verifiable and portable form that can be consumed by on-chain systems. Compliance checks such as Know Your Customer (KYC) verification, sanctions screening, or transaction monitoring are conducted within existing institutional frameworks, with their outcomes encoded as verifiable attestations or credentials that can be attached to transactions and evaluated programmatically.

Project Mandala led by the Bank for International Settlements[6] exemplifies this approach. It demonstrates how jurisdiction-specific compliance requirements can be codified into machine-readable rules and expressed as verifiable artefacts. These accompany digital settlement flows and can be evaluated within transaction execution logic without requiring disclosure of the underlying data. This supports automation, interoperability, and alignment with regulatory expectations in cross-border contexts.

Current implementations of this pattern typically operate within defined institutional trust boundaries, where participating entities agree on the standards, formats, and validity of compliance representations.

Box 1: Example – Project Mandala

The BIS Innovation Hub's Project Mandala explores a "compliance-by-design" approach to improve the efficiency and robustness of compliance verification for cross-border payments and digital asset transactions in a privacy-preserving manner. It enables jurisdiction-specific compliance requirements to be incorporated directly into transaction workflows at the pre-validation stage of the transaction and provide strong compliance execution assurance while maintaining the desired data privacy for sensitive information. The project contributes to central banks' mandate of maintaining the safety, security and integrity of payments market infrastructures.

In Phase 1, the project demonstrated the feasibility of automating compliance procedures through pre-validation of regulatory requirements and the use of privacy-preserving proofs of compliance for sanctions screening and capital flow management use cases. It also demonstrated interoperability of this approach in traditional and tokenised financial systems. Currently in Phase 2, which is led by the BIS Innovation Hub Singapore Centre with Bangko Sentral ng Pilipinas, Banque de France, Bank Negara Malaysia, Central Bank of Kuwait, Reserve Bank of Australia, Reserve Bank of India, and the Monetary Authority of Singapore, Project Mandala is expanding its scope to empirically

[6] The first phase of Project Mandala was conducted with the Reserve Bank of Australia, Bank of Korea, Bank Negara Malaysia, and the Monetary Authority of Singapore. The proof of concept explored the use of cryptographic techniques, including zero-knowledge proofs and secure multi-party computation, to support privacy-preserving verification of compliance outcomes (Bank for International Settlements Innovation Hub, 2024).

evaluate compliance pre-validation against existing processes and to cover a broader set of digital assets and payment ecosystems.

Mandala's key features include:

- Configurable pre-validation compliance controls that enable financial institutions to define and apply jurisdiction-specific compliance checks across different transaction categories before payment execution.
- Integration with existing compliance technology stacks, allowing Mandala to work with existing compliance software without requiring significant infrastructure replacement.
- Modular architecture to leverage different privacy-preserving and verifiable computation approaches to generate cryptographically verifiable proofs of compliance.
- Privacy-preserving information sharing mechanisms that enable participants to validate regulatory compliance and transaction eligibility while minimizing unnecessary disclosure of sensitive customer or transaction data.

Mandala high-level architecture

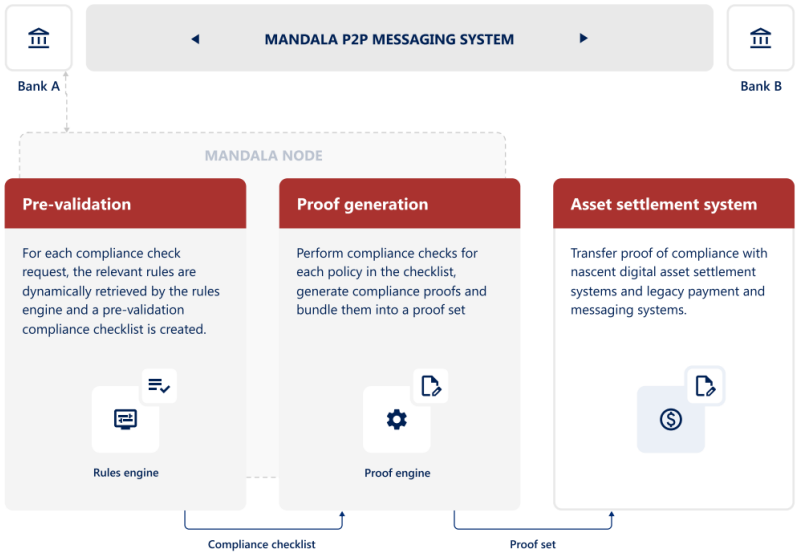

Source: Project Mandala Phase 1

Generation of verifiable compliance proof in Mandala can serve as trusted input into programmable compliance smart contracts, enabling automated policy enforcement, conditional transaction execution, and auditable compliance verification across cross-border payment and digital asset workflows.

Approach 3: Verifiable Identity as a Foundation for Participation

A third approach focuses on verifiable identity infrastructures, which underpin the ability to determine who is authorised to interact with regulated digital assets.

Identity is represented through portable, cryptographically verifiable credentials that assert the legal status, roles, and attributes of participating entities. The Global Legal Entity Identifier Foundation (GLEIF) has advanced this through standards for verifiable Legal Entity Identifiers (vLEIs), enabling organisations to assert their

identity in a manner that can be independently verified[7]. Complementary efforts by IDEMIA explore how real-world identity verification processes, including KYC and biometric identity proofing, can be operationalised through decentralised identifiers and verifiable credentials for digital asset environments[8].

These identity layers bind on-chain activity to real-world entities, ensuring that only authorised and eligible participants can hold, transfer, or administer regulated tokens, and provides a bridge between distributed infrastructures and existing regulatory regimes.

Approach 4: Modular and Composable Compliance Models

A further evolution is the emergence of modular and composable models, where compliance functions are separated into discrete components that can be combined across different participants and transaction contexts. These participants may include issuers, intermediaries, identity providers, verification services, custodians, and infrastructure operators.

Through interoperable components, these models allow identity management, policy enforcement, monitoring, and portability to be coordinated across tokenised asset systems. For example, Chainlink's Automated Compliance Engine (ACE) is described as a modular framework that supports identity management, policy enforcement, monitoring, and cross-chain portability for compliance-focused digital assets[9]. Token standards such as ERC-3643 (or T-REX protocol), implemented by platforms like Tokeny, embed compliance rules directly into token contracts while allowing external identity and verification providers to supply required inputs[10].

This modular approach supports flexibility and scalability, particularly in cross-border contexts involving multiple jurisdictions and service providers. It also requires clear allocation of responsibility across components, including how external providers are selected, monitored, updated, and relied upon[11].

---

[7] The vLEI framework extends the LEI into a verifiable credential model, supporting digitally verifiable organisational identity and related role assertions across systems (Global Legal Entity Identifier Foundation, 2024).

[8] IDEMIA frames decentralised identifiers and verifiable credentials as mechanisms for privacy-preserving and cross-jurisdictional KYC in digital asset environments, enabling users to share necessary identity attributes through cryptographically signed credentials while supporting biometric binding for higher assurance (IDEMIA, 2025).

[9] ACE is framed as a modular compliance stack, illustrating how identity, policy, monitoring, and interoperability functions can be coordinated across networks rather than embedded entirely within a single token contract (Chainlink, 2025).

[10] The T-REX approach is designed around modular token-level restrictions linked to identity and claim verification, allowing issuers to retain control over transfer eligibility without embedding all compliance data in the token contract itself (Tokeny Solutions, 2023).

[11] This is consistent with financial-sector expectations that regulated institutions remain responsible for managing risks arising from reliance on external service providers, including through due

Cross-Cutting Considerations

Across these approaches, the direction of travel is similar. Programmable compliance is moving towards systems that are more verifiable, privacy-preserving, and modular, with a clearer separation between data, logic, and execution.

Most current implementations remain institutionally bounded, relying on predefined governance arrangements, including trusted issuers of identity credentials, recognised validators of compliance representations, and agreed mechanisms for updating policy rules.

At the same time, broader market developments point toward increasing interaction with public or open ledger networks, including networks operated outside a single institution's organisational boundary. This is driven by their liquidity, composability, and reduced need for institution-specific infrastructure. Future programmable compliance frameworks may therefore need to operate across hybrid environments, where permissioned features, such as controlled access, approved participants, and enforceable asset controls, interact with open-network features, such as shared liquidity, composability, and broader connectivity. This requires mechanisms that can maintain regulatory assurances while allowing tokenised assets and financial applications to interoperate across different ledger environments.

---

diligence, governance, ongoing monitoring, audit rights, business continuity arrangements, and exit management (Monetary Authority of Singapore, 2026).

GL1 Programmable Compliance

Compliance obligations need to keep pace with evolving threat vectors and risk landscapes, while core transaction and asset logic need to operate with stability, predictability, and assurance. Smart contract logic is typically difficult to change once deployed, and updates require formal upgrade mechanisms and third-party assurance. In regulated tokenised asset environments, compliance is also organised across multiple controls and points of responsibility, meaning that policy logic, decision evidence, and execution pathways are not always aligned. A structured approach that keeps policy rules separate from transaction and asset execution logic is therefore preferred, so that changes to one do not require changes to the other.

Table 1: STAR Framework for Compliance Requirements

Pillar	Focus	Typical compliance activities (illustrative)
Status (S)	Participant eligibility	Customer due diligence (KYC/KYB), sanctions screening, beneficial ownership where relevant, investor eligibility checks, jurisdictional access controls[17].
Transaction (T)	Flow and movement	Risk-based thresholds and velocity limits, Travel Rule information exchange, counterparty and corridor-based restrictions, transaction holds or conditional routing where required checks are not met.
Asset (A)	Instrument state	Corporate actions and administrative actions (for example, redemption, issuance, supply controls)[18], transfer restrictions embedded in the instrument, interventions such as freezes or forced transfers where legally supported and governed[19].
Reporting (R)	Data and auditability	Suspicious activity escalation triggers and evidence packaging, tax reporting outputs (for example, broker reporting), cap table snapshots where applicable, reserve disclosures or attestations (where used) and audit trails[20].

The model presented in this section is organised around this principle using the STAR framework, comprising of *Status*, *Transaction*, *Asset*, and *Reporting*. The framework classifies compliance requirements according to whether they relate to

[17] FATF standards treat customer due diligence, beneficial ownership, and sanctions-related controls as core elements of participant-level compliance.
[18] Project Guardian's fixed income framework illustrates token lifecycle functions and related administrative actions in tokenised asset settings (Monetary Authority of Singapore, 2023b).
[19] FATF's asset recovery guidance provides a basis for interventions such as freezes, restrictions, and transfers where these are legally supported and subject to proper authority.
[20] International Organization of Securities Commissions' (IOSCO) assessment methodology highlights the importance of records, disclosure, and supervisory review in regulated markets. The examples listed here reflect reporting and auditability functions that support oversight and accountability (IOSCO, 2011).

participant eligibility, transaction permissibility, asset-level controls, or evidencing outputs. This allows the architecture to avoid collapsing different compliance functions into a single monolithic layer of execution logic. Organising compliance controls in this way supports clearer delineation of responsibilities, permissioning, and assurance boundaries, and reduces the likelihood that changes in one domain affect controls addressing other domains.

The following subsections describe the key features of the model and how they interact to support programmable compliance across the STAR pillars.

Key Features of GL1 Programmable Compliance Model

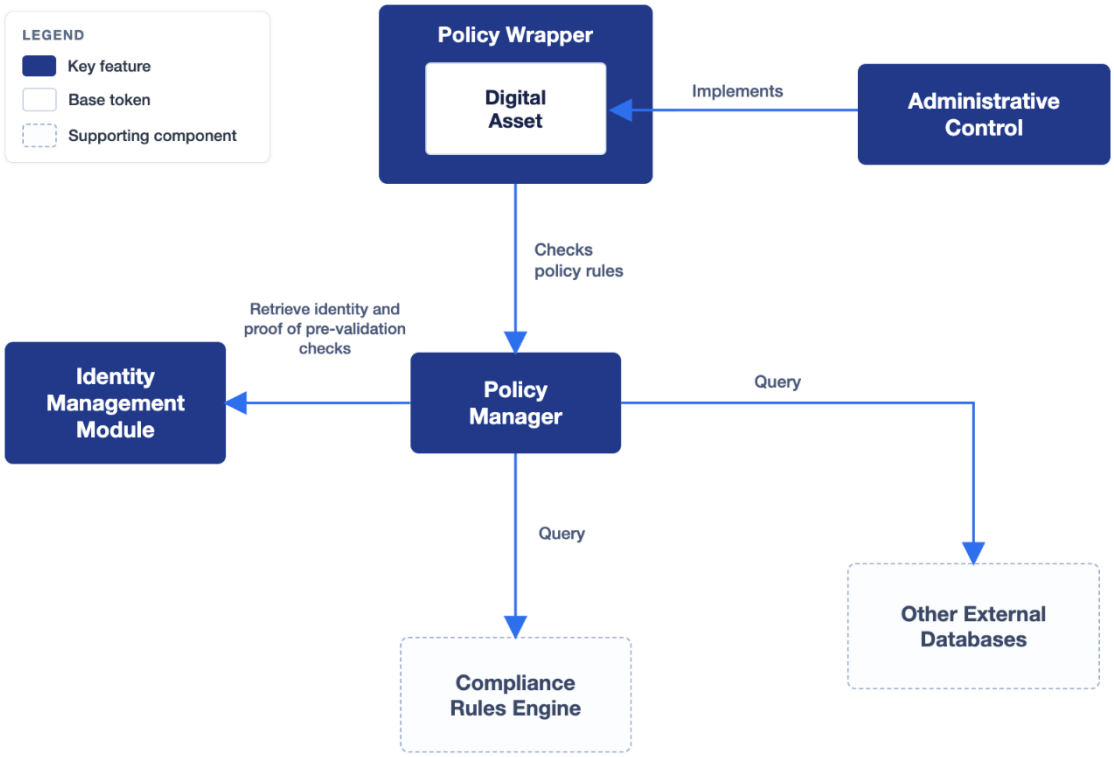

Figure 1: GL1 Programmable Compliance Model

Identity Management

The Identity Management feature supports the *Status* pillar by providing the entity-level view needed for financial compliance. In tokenised asset environments, a single individual or institution may transact through multiple wallet addresses, accounts, custodians, or networks, while compliance obligations often attach to the underlying entity rather than to a specific wallet address. The Identity Management Module closes this gap by linking relevant addresses and identifiers to the same underlying entity.

The module may support different levels of implementation. At a basic level, it may enable address whitelisting for participants that have completed required checks. At a more advanced level, it may link multiple addresses to a single entity through

attestations, cryptographic proofs, or integration with existing customer identification systems. It may also support credentials issued by trusted third parties, including credentials relating to Know Your Customer (KYC), sanctions screening, beneficial ownership, investor eligibility, or other compliance attributes.

The Identity Management Module does not need to perform primary identity verification itself. Its role is to maintain the trust criteria for accepting identity-related credentials, resolve entity references, and support revocation, expiry, and refresh of Status evidence over time. Jurisdictional differences in KYC standards, eligibility criteria, and sanctions regimes also mean that *Status* evidence may need to be scoped to specific jurisdictions or transaction contexts, rather than treated as universally portable.

Box 2: Example - Verifiable Credentials with GLEIF

Verifiable credentials are cryptographically signed digital proofs that can allow transaction systems to verify identity, authority, and compliance attributes without repeatedly exchanging underlying documents. Their usefulness depends on both the integrity of the credential issuance process and the reliability of the real-world identifier to which the credential refers.

GLEIF's vLEI links digital credentials to the Legal Entity Identifier (LEI) system, providing a trusted organisational identity credential for digital asset ecosystems. When attached to digital ledgers, assets, or transaction workflows, vLEIs can help verify legal entities (wallet holders, token issuers, reserve attestation issuers, etc.), authorised representatives, and associated compliance attestations such as KYC, AML, sanctions, or investor eligibility checks. This supports more automated, interoperable, and machine-readable transaction compliance across tokenised asset ecosystems.

Representation of a vLEI credential content

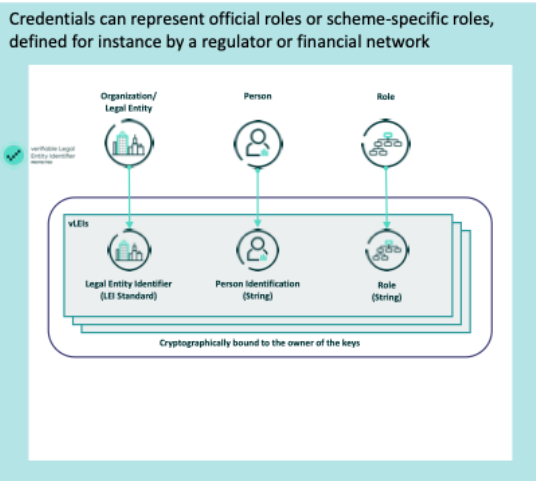

The root of trust architecture

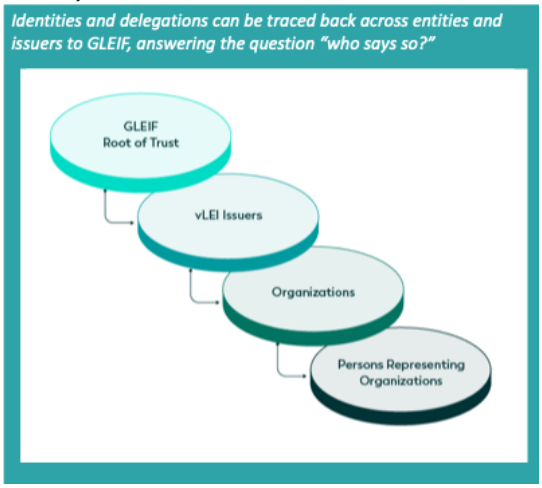

Policy Wrapper

The Policy Wrapper supports the *Transaction* pillar by providing the execution boundary for compliant transfers. Drawing on the concept of Purpose Bound Money

(PBM)[21], the Policy Wrapper separates the underlying asset from the conditions governing how it may be used or transferred. Rather than embedding compliance logic directly into the core token contract, the wrapper locks the underlying base asset token and issues a corresponding wrapped token. Transfers then take place through the wrapped token, while the base asset remains locked within the wrapper.

Box 3: How PBM Works

PBM, formalised in ERC-7291 Purpose Bound Money, demonstrates how digital assets can be structured such that their use, transfer, and ultimate redemption are governed by predefined conditions. Rather than modifying the underlying asset, PBM introduces a wrapping mechanism that binds the asset to a defined "purpose of payment", ensuring that value can only circulate within permitted contexts.

PBM adopts a lock-and-mint model, where an underlying token is locked and a purpose-bound representation is issued. This establishes the initial binding of policy conditions to the asset. However, the significance of PBM lies not only at issuance, but in how these conditions persist across the asset's lifecycle, governing both subsequent transfers and eventual redemption.

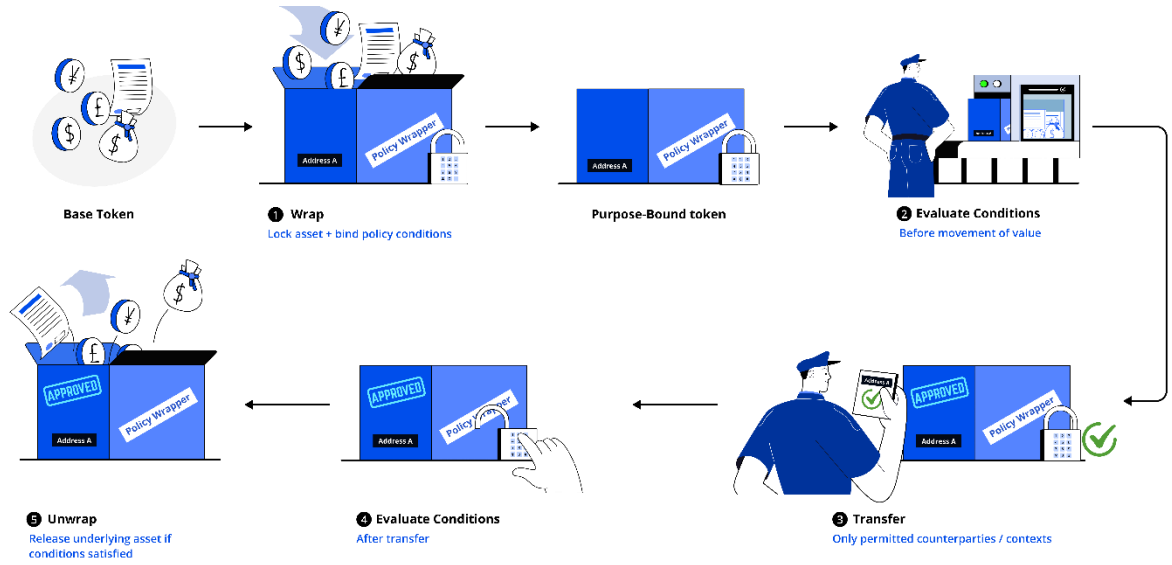

PBM enforces policy conditions across three stages:

- **At wrapping**, conditions are bound to the asset before it enters circulation, ensuring that only purpose-constrained representations can be used.
- **During transfer**, the wrapped token can circulate, but only within permitted boundaries. Transfer logic can incorporate conditions such as restrictions on counterparties or usage contexts, allowing compliance checks to occur *prior to value exchange*. This enables control not just at issuance, but throughout secondary transactions.

[21] PBM is an architectural approach in which an underlying digital asset is paired with a wrapper that enforces conditions on transfer or use (Monetary Authority of Singapore, 2023a). Its relevance here lies in the separation between the base asset and the conditional execution layer, allowing policy conditions to be applied without changing the core token contract.

- **At unwrapping**, the underlying asset is released only when specified conditions are met. If conditions are not satisfied, the asset remains locked, preventing non-compliant actors from accessing the underlying value.

PBM demonstrates that wrapping is not a one-time control but a persistent enforcement layer. Assets are bound to conditions at entry, constrained throughout circulation, and conditionally released at exit, ensuring that compliance is enforced continuously across the entire transaction lifecycle. The model ensures that compliance is enforced continuously across the transaction lifecycle.

When a transfer of the wrapped token is attempted, the wrapper initiates the compliance evaluation process. It checks whether the relevant policy conditions have been satisfied before allowing the transaction to proceed. This allows compliance conditions to be applied at the transfer layer without modifying the underlying asset contract.

This structure allows the same base asset token to operate under different compliance conditions. Different wrappers may be used to reflect jurisdictional requirements, market-specific rules, or policy changes over time. Where unwrapping is permitted, the wrapper performs the required checks before releasing the locked base asset token.

The Policy Wrapper also supports the *Reporting* pillar by receiving and, where required, recording, emitting, or relaying the Compliance Attestation generated through the evaluation process (discussed in the following subsection). In this sense, the wrapper does not only determine whether a transaction can proceed; it also provides the point at which transaction-level compliance evidence can be preserved for audit, supervisory review, or regulatory reporting.

Policy Manager

The Policy Manager supports the *Transaction* pillar by managing the policies that define which compliance rules apply to a given transaction. These policies may specify the relevant rule set, required Status evidence, transaction thresholds, jurisdictional conditions, counterparty requirements, reporting triggers, escalation rules, and precedence arrangements where multiple rules apply.

The Policy Manager determines how transaction context affects the applicable compliance requirements. For example, a policy may require different checks depending on value, velocity, corridor, counterparty type, asset type, or delivery channel. It may also specify when participant-level Status evidence is required to satisfy a transaction-level condition.

The Policy Manager separates policy management from transaction enforcement. The wrapper requests an evaluation through a standardised interface, while the Policy Manager identifies the applicable policy and determines which rules need to be checked. It may then coordinate with the Identity Management Module, external

compliance services, or other supporting modules to obtain the required inputs and evaluation results.

Where multiple rules or modules produce inconsistent outputs, the Policy Manager should apply defined governance rules on precedence, timeout handling, escalation, and fallback treatment. These arrangements are not purely technical. They reflect institutional risk appetite, regulatory obligations, and accountability arrangements, and should therefore be subject to governance review.

Box 4: Example - Chainlink Automated Compliance Engine

The Chainlink Automated Compliance Engine (ACE) is a modular framework for enforcing compliance policies on any blockchain or digital asset the moment a transaction is executed. Programmable compliance rules live separately from application logic, allowing issuers and supervisors to add, layer, update, or remove policies as regulations evolve—without needing to redeploy the underlying contract.

ACE is an example of how the GL1 Programmable Compliance Model can be implemented in practice, featuring a Policy Manager that orchestrates the Identity Management and Compliance Rules Engine layers, which the Policy Wrapper may invoke to enforce compliance on every transaction.

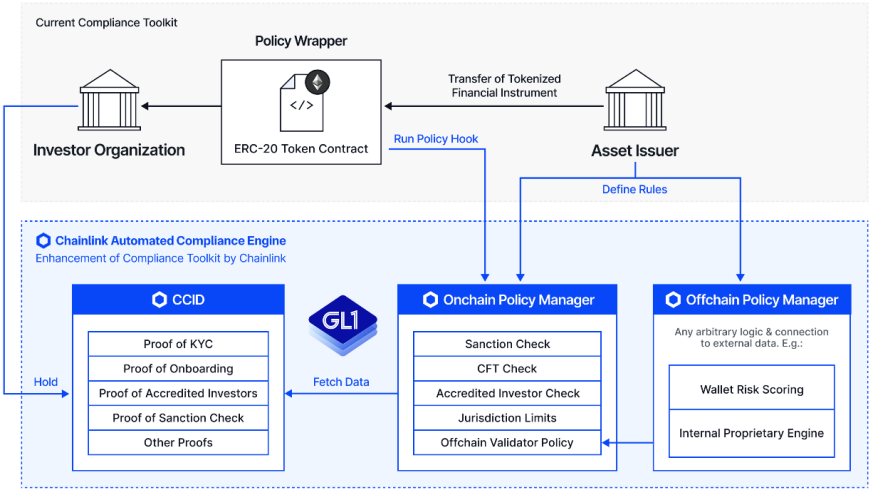

Policy Engine and Composable Policies: The Policy Engine is an on-chain contract that evaluates an ordered set of modular policies whenever a protected operation—such as transfer, mint, burn, lock—is invoked. Each policy returns: Allow, Reject, or Continue, allowing institutions to assemble compliance logic from reusable building blocks, including allow lists, sanctions screening, jurisdictional limits, accredited-investor checks, transaction thresholds, identity-gated transfers, and more. Policies can be layered on a single operation, reused across operations, and swapped as regulations evolve, without redeploying the protected contract. Where the same asset must operate under multiple jurisdictional regimes, separate Policy Wrappers can be deployed for the same underlying asset, each connected to a Policy Engine and its own policy chain.

Compliance computation that does not fit on-chain—such as complex risk scoring or checks from internal proprietary systems—is performed off-chain and attested via cryptographically-signed results, which an on-chain validator policy verifies before the transaction proceeds.

Every policy evaluation emits an on-chain event, providing a complete audit trail.

Cross-Chain Identity (CCID): CCID is a portable identity layer. A CCID links a holder's wallet addresses across networks to a single identity. Credentials such as proof of KYC, sanctions screening, or accredited-investor status are issued once against the CCID, and recognised on every network on which CCID is supported, eliminating duplicative onboarding when the holder transacts on a new network.

Some policies require computation before a compliance decision can be reached. For example, a policy may require the system to calculate cumulative transaction values, apply risk scores, or compare counterparties against sanctions lists. The Compliance Rules Engine supports the Policy Manager by performing these computations and returning whether the relevant policy check has passed, failed, or requires further handling.

The Compliance Rules Engine may be implemented within or outside of tokenised asset environments, or in hybrid form. On-chain execution can provide transparency and deterministic auditability, but may be costly or inflexible for complex rules and large datasets. Off-chain execution can support more sophisticated computation and easier integration with external data sources, but requires controls over integrity, availability, audit trails, and data provenance. A hybrid approach may place simpler checks on-chain while reserving more complex analysis for off-chain systems.

In this architecture, the Compliance Rules Engine does not determine which policy applies and does not enforce the transaction outcome directly. The Policy Manager identifies the applicable policy and required checks, the Compliance Rules Engine evaluates whether those checks are satisfied, and the Policy Wrapper enforces the resulting outcome at the transaction boundary.

Administrative Control

Administrative Control supports the *Asset* pillar by providing intervention capabilities over the tokenised asset itself. In regulated settings, institutions may need to restrict addresses, freeze assets, recover assets, force transfers, or trigger emergency measures in response to suspicious activity, legal directives, operational incidents, or broader risk events. Without such controls, tokenised assets may remain technically transferable even where an institution is required to intervene.

These controls operate at the asset layer rather than the ordinary transaction flow. They may include address-level restrictions, recovery functions that are distinct from user-initiated transfers, and mechanisms to pause or limit regular transfers while preserving necessary administrative access. This allows the architecture to act directly on the asset where ordinary transfer controls are insufficient or where intervention is required after issuance.

Administrative Control should therefore be understood as a separate intervention capability within the model. It does not determine participant eligibility, define transaction policies, or evaluate whether a policy condition has been satisfied. Its function is to provide bounded mechanisms for acting on the tokenised asset itself when required.

How the Features Work Together

These features work in concert through a coordinated transaction flow built around a central concept: the Transaction Envelope. This a structured record that travels with each transaction, accumulating context and approvals as it moves through the system, ensuring that by the time a transaction is executed, it has been enriched with relevant data, evaluated against defined rules and attested by the appropriate parties.

When a transaction is initiated, the Policy Wrapper constructs the Transaction Envelope. The envelope captures core transaction details and may include data such as asset identifier, transaction amount, transaction datetime, and a set of Party Packets representing the actors involved. These Party Packets are designed to accommodate different mediation patterns, including Entity-to-Entity, Entity-to-Virtual Asset Service Provider (VASP), VASP-to-Entity, and VASP-to-VASP interactions, and may contain information such as role, wallet address, jurisdiction, and pre-validation proofs. These data elements are illustrative rather than prescriptive, and the precise structure of the Transaction Envelope is expected to evolve as standards develop.

At the point of construction, Party Packets contain only minimal identifiers required to reference the relevant parties. This supports the separation of identity data from transaction execution. The Identity Manager is subsequently used to enrich the Transaction Envelope by resolving these identifiers into the required identity attributes and attestations for compliance evaluation, reducing the exposure of sensitive information within the transaction layer while ensuring that sufficient data is available for downstream checks.

Figure 2: Illustration of the Transaction Envelope and Party Packets

The enriched Transaction Envelope is then submitted to the Policy Manager. The Policy Manager identifies the policies relevant to the transaction and passes the Transaction Envelope to those policies for evaluation. Where a policy requires computation to reach a decision, such as calculating transaction thresholds, checking cumulative exposure, applying risk scores, or comparing data against external lists, the Compliance Rules Engine supports that policy by performing the required computation.

The results of the applicable policy evaluations are consolidated by the Policy Manager into a Compliance Attestation[22]. The Compliance Attestation represents the overall compliance outcome of the transaction and may include structured results corresponding to individual checks. It is then returned to the Policy Wrapper, which applies the outcome at the point of execution. Transactions that satisfy the required conditions are permitted to proceed, while transactions that do not satisfy those conditions are rejected or handled according to the applicable control logic.

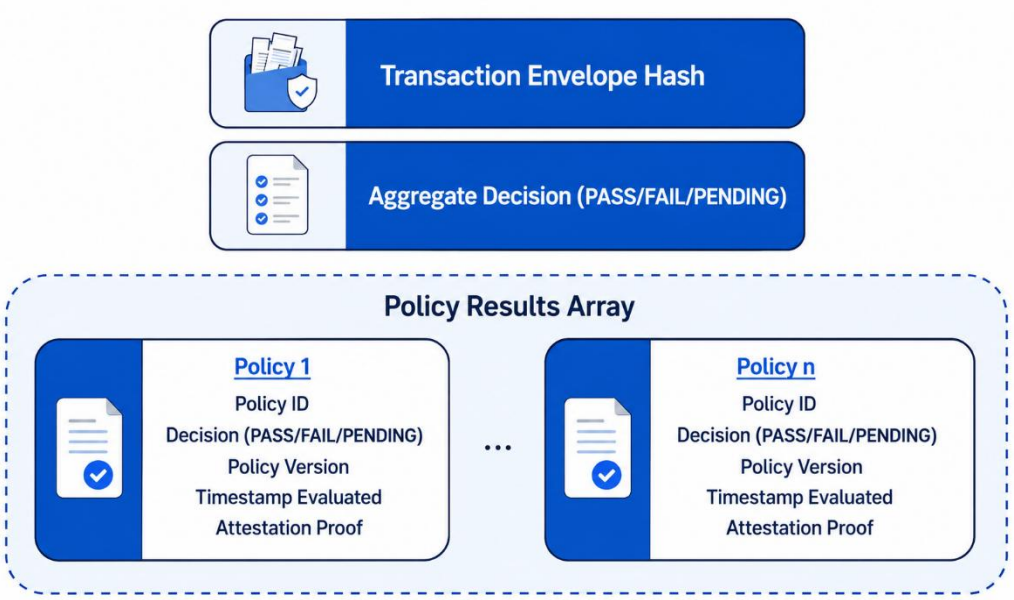

Figure 3: Compliance Attestation Data Structure

Administrative Control operates alongside this transaction flow as an asset-level intervention capability. It is not necessarily invoked for every transaction. Rather, it provides the mechanism through which the institution may act directly on the

[22] The Compliance Attestation is an immutable execution-time record: it captures the compliance evaluation outcome for a specific transaction at the moment of evaluation, against the policy version and trust dependencies active at that time, and does not carry forward as a general authorisation to transact. It is distinct from the pre-validation credentials and identity attestations that may be consumed as inputs—those are issued by trusted third parties through the Identity Manager and are subject to their own expiry, revocation, and refresh lifecycle. The integrity and authenticity of the Compliance Attestation, including any cryptographic binding, are implementation choices that should be governed in accordance with the evidencing requirements set out in the Governance and Accountability section.

tokenised asset where required, such as by restricting transfers, freezing assets, or recovering assets. The Transaction Envelope and Compliance Attestation may provide relevant evidence for such intervention, but the intervention itself is carried out through Administrative Control rather than through the ordinary wrapped-token transfer flow.

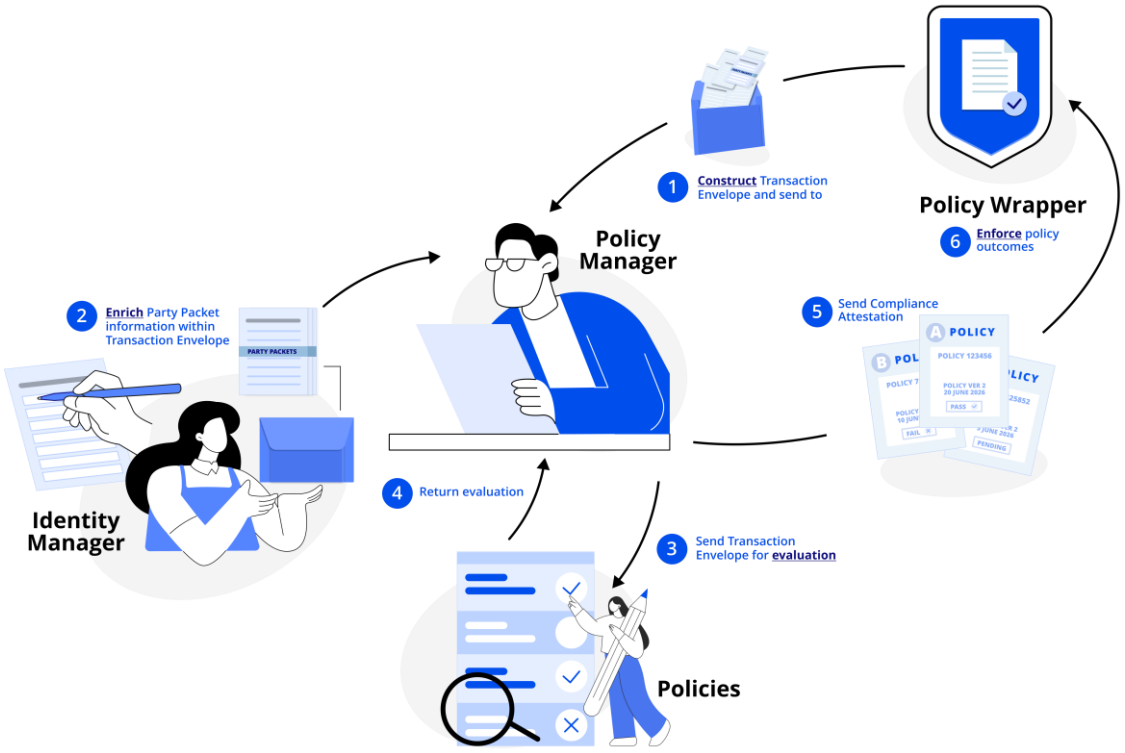

Figure 4: Transaction Envelope flow: Construction, Enrichment, Evaluation, and Enforcement

This transaction flow maintains clear boundaries between execution, identity management, and compliance evaluation, while preserving an integrated end-to-end process. It aligns with the structured view of compliance requirements set out earlier, with *Status*-related requirements incorporated through Identity Manager enrichment, *Transaction*-level requirements evaluated through the Compliance Rules Engine, *Asset*-level enforcement applied through the Policy Wrapper, and administrative controls available where intervention is required. *Reporting*-domain obligations are addressed through the Compliance Rules Engine, which evaluates reporting triggers and generates compliance signals used to satisfy downstream reporting requirements; the evidencing records described in the Governance and Accountability section serve as the primary vehicle for audit trail and supervisory reporting purposes.

By structuring compliance around the Transaction Envelope and associated attestations, the model supports consistent enforcement, modular system design, and adaptability to evolving regulatory requirements[23].

Risk and Governance Considerations

Programmable compliance changes how compliance controls are implemented, evaluated, and evidenced within regulated financial activity. Embedding compliance logic closer to transaction execution can improve consistency, reduce reconciliation overhead, and support more structured evidentiary outputs across tokenised asset environments. At the same time, this architectural shift introduces new considerations around how compliance risk is managed and governed.

Some risks remain present even with programmable execution. Sanctions data may become outdated, identity information may not propagate in time, reporting obligations may remain fragmented across jurisdictions, and legal uncertainty may persist in cross-border activity. Programmable compliance can reduce some coordination and timing frictions, but it also increases dependency on automated decisioning, external data integrity, smart contract design, administrative intervention mechanisms, and interoperability between system components.

Additional risks arise where compliance controls are embedded within transaction workflows, policy management layers, and asset-level controls. Compliance outcomes may depend on oracle reliability, credential propagation, policy configuration, administrative override mechanisms, and interactions between multiple compliance domains operating at the same time. In tokenised asset environments with atomic or near-real-time settlement, failures may also propagate with limited opportunity for manual review or post-settlement remediation.

This section examines these residual and emergent risks, and the governance considerations needed to manage them.

---

[23] The architecture's component structure maps to key AML/CFT obligations at the functional level. As illustrative examples: Travel Rule requirements (FATF R.16) are supported through the Transaction Envelope and Identity Manager enrichment, with execution conditioned through the Policy Wrapper; suspicious activity reporting obligations (FATF R.20) are supported through the Compliance Rules Engine, which evaluates suspicion triggers and preserves evaluation context for both executed and rejected transactions in evidencing records; and sanctions screening obligations (FATF R.6–7) are supported through the Identity Manager and enforced at execution through the Policy Wrapper, with asset-level intervention capability held by Administrative Control. IOSCO Principles relating to investor protection and market conduct similarly inform the design of eligibility controls and transaction monitoring. These mappings are functional and illustrative; whether the architecture satisfies applicable obligations in any specific jurisdiction depends on the relevant regulatory framework.

Residual Risks

Temporal Displacement and Stale Compliance States

Programmable compliance systems continue to depend on the timelines and accuracy of external information. Sanctions status, beneficial ownership information, insolvency events, investor eligibility, and counterparty standing may change after an attestation has been issued or after a transaction has been initiated. Where settlement occurs atomically or near real-time, the operational window available for re-evaluation may narrow significantly.

This risk is particularly acute in long-dated or multi-stage transactions such as securities financing arrangements, collateral substitution, and deferred settlement obligations. Architectures should therefore distinguish between point-in-time attestations and continuously valid attestations, and should define explicit re-evaluation triggers for transactions that remain open over extended periods.

Cross-Jurisdiction Alignment

Cross-border financial activity routinely engages multiple legal and supervisory regimes simultaneously. Comparable compliance objectives may be implemented differently across jurisdictions, including differences in reporting thresholds, investor eligibility rules, sanctions handling, Travel Rule implementation, and administrative intervention requirements[24].

Programmable compliance architectures may reduce duplication in evidencing and transaction processing, but they do not remove the need to evaluate jurisdiction-specific obligations. Policy Managers operating across multiple jurisdictions may therefore encounter conflicting requirements or incompatible execution conditions. Where these conditions are not resolved through explicit coordination logic, transaction deadlocks, inconsistent compliance outcomes, or unintended permissive execution states may arise.

Reporting Compatibility

Compliance attestations generated within programmable environments do not automatically correspond to supervisory reporting artefacts required under existing regulatory frameworks. Suspicious activity reports, trade reporting obligations, collateral disclosures, and prudential reporting requirements continue to depend on

[24] A BIS-led roundtable convened during the Singapore FinTech Festival noted that, although regulations are often based on similar principles, financial institutions face significant compliance challenges due to differences in how these rules are applied across jurisdictions (Global Finance & Technology Network, 2023).

formats and evidentiary standards designed primarily for institutional reporting environments[25].

Programmable compliance architectures may improve the consistency and traceability of compliance evidence internally, while still requiring transformation, aggregation, or interpretation before outputs become usable for supervisory or investigative purposes. Reporting systems should therefore preserve human-readable evidentiary records and maintain compatibility with supervisory reporting frameworks operating outside tokenised environments.

Novel Risks with Programmable Compliance

Smart Contract Rigidity

On-chain compliance logic executes deterministically. While this supports consistency and predictability, it may also create situations where legally or operationally appropriate exceptions cannot be accommodated without governed override mechanisms.

Incorrectly parameterised compliance logic may therefore produce outcomes that are technically valid but operationally inappropriate or legally disputed. In tokenised settlement environments, these effects may become difficult to reverse once settlement finality has been reached.

Oracle Dependency

Programmable compliance systems depend extensively on external information sources, including sanctions databases, identity registries, market data feeds, jurisdictional risk indicators, and valuation services. Oracle infrastructure bridges these external systems with programmable execution environments.

Failures within oracle infrastructure may therefore propagate directly into transaction outcomes. Delayed sanctions updates, inaccurate valuation feeds, unavailable identity services, or corrupted external data may produce non-compliant execution outcomes despite technically correct internal processing. Centralised oracle

[25] For example, FATF Recommendation 16 requires originating VASPs to collect and transmit structured originator and beneficiary information—to be formatted in accordance with applicable standards such as IVMS101—for transfers above applicable thresholds. The use of zero-knowledge proofs or encrypted references as a mechanism to support this transmission while limiting on-chain Personally Identifiable Information exposure is an active area of technical and regulatory development. Several Travel Rule protocol providers have developed off-chain communication frameworks that decouple the on-chain settlement event from the transmission of structured originator/beneficiary data; whether any specific approach satisfies applicable requirements depends on the jurisdiction. FATF is also actively updating its guidance for virtual asset environments. Institutions should confirm the permissible implementation approach with the relevant regulatory authority.

architectures may also create concentration risks where the failure of a single dependency affects multiple institutions or transaction flows simultaneously.

Privacy, Immutability, and Data Disclosure

Programmable compliance architectures operating on tokenised asset infrastructure may increase exposure to privacy and data governance risks, particularly where transaction metadata, identity-linked references, or compliance evidence become visible across shared or persistent infrastructure layers.

In traditional financial systems, customer and transaction information is typically held across institutions, intermediaries, and reporting channels. In tokenised environments, transaction traceability, interoperability requirements, and evidentiary consistency may increase pressure for more structured and portable compliance data. Where personally identifiable information or recoverable identity references become associated with immutable transaction records, legal obligations relating to confidentiality, data minimisation, rectification, or erasure may become operationally difficult to satisfy.

Privacy-enhancing approaches such as selective disclosure mechanisms, encrypted references, and zero-knowledge proofs may reduce direct exposure of sensitive information, but they also introduce additional dependencies relating to credential management, metadata correlation, and evidentiary retrieval by competent authorities.

Box 5: Example - Bermuda Protocol[27] privacy solution

Bermuda is a set of permissionless smart contracts, compatible with standard signing infrastructure, deployable on Ethereum Virtual Machine-compatible networks. It uses client-side zero-knowledge proofs to enable private transactions while enforcing asset- and transaction-level compliance policies on every operation.

Once an asset is deposited within the Bermuda system, participants can transfer it privately, execute repo trades and swaps, or interact with third-party decentralised finance protocols — all without exposing transaction details on-chain. An operation is only finalised once all applicable policies are satisfied. The lifecycle concludes with a withdrawal that releases the asset to the beneficiary's destination wallet.

[27] For avoidance of doubt, ‘Bermuda’ in this document refers solely to the privacy protocol and shall not be construed as a reference to the British Overseas Territory of Bermuda.

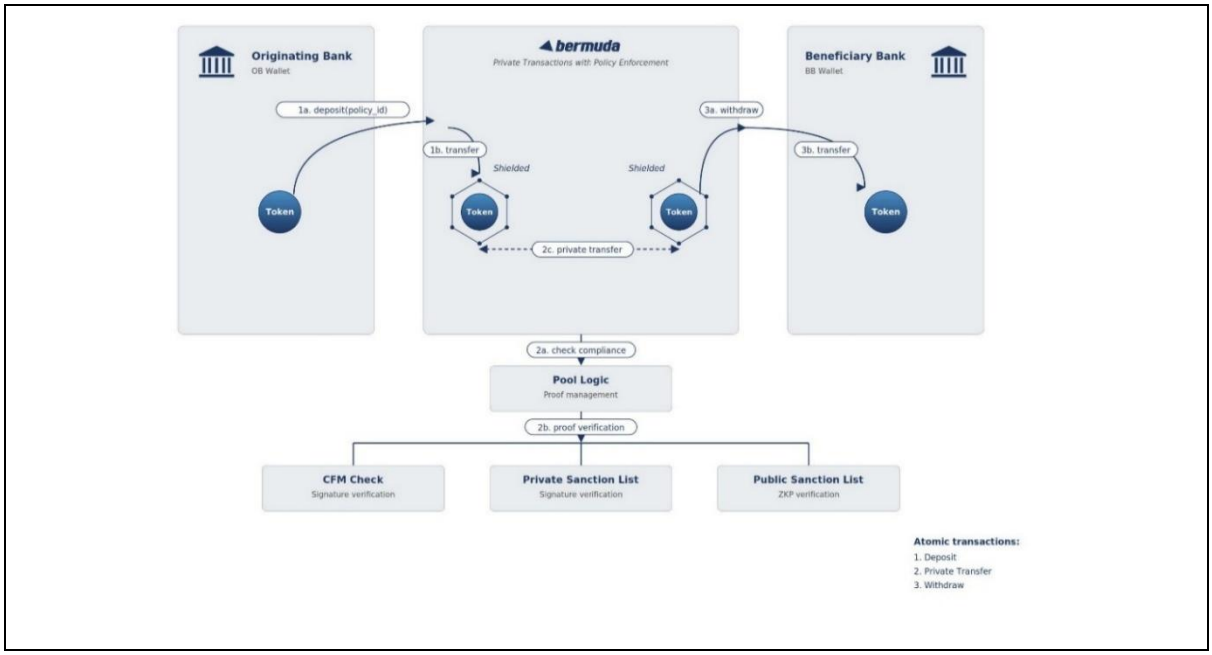

Administrative Intervention and Governance Concentration

Certain tokenised asset standards contemplate administrative capabilities such as freezing assets, reversing transactions, recovering funds, or enforcing court orders. These functions may be operationally necessary in regulated environments, particularly for settlement finality disputes, sanctions enforcement, fraud remediation, or insolvency administration.

However, these capabilities also create governance concentration risks. Administrative keys or override authorities may become high-value attack targets or sources of insider abuse. Governed intervention mechanisms should therefore incorporate explicit role separation, approval controls, audit logging, escalation pathways, and evidentiary retention requirements.

Interface and Coordination Failure

Implementations of the Programmable Compliance model defined in this paper will involve interactions between multiple components, including Policy Wrappers, Policy Managers, Identity Managers, and oracle services. Failures at these interfaces may produce outcomes that do not reflect the intended policy state.

Examples include degraded connectivity between identity services and transaction engines, inconsistent policy propagation across jurisdictions, malformed evidentiary outputs, or sequencing failures between settlement and compliance validation layers. Architectures should therefore define default behaviours at component boundaries and governed fallback procedures where policy dependencies cannot be resolved automatically.

Governance and Accountability

Governance should distinguish between the source of a requirement and the architectural object through which that requirement is operationalised. Legal, regulatory, and supervisory obligations provide one source of requirements. Institution-defined controls, such as risk appetite, product parameters, and operating rules, provide another. These requirements become operational only when translated into governed objects within the architecture, such as policy artefacts, trust dependencies, execution conditions, administrative capabilities, and evidencing records.

Figure 5: Governance application to the artefacts

This distinction matters because the architecture separates relatively stable components from more changeable policy and control objects. The base asset token and execution surfaces may remain stable, while policy parameters, trusted attestations, eligibility criteria, and reporting triggers may change over time. Governance should therefore ensure that such changes are attributable, reviewable, and traceable without requiring unnecessary redesign of the enforcement layer.

A related principle is that accountability should follow control over the relevant governed object. In the model, the issuer of the base asset, the operator of the Policy Wrapper, the manager of policy artefacts, attestation or verification providers, technology providers, and participants may each control different parts of the compliance architecture. Accountability should therefore be allocated according to the function each party controls and the extent to which that function affects compliance outcomes, evidencing, or asset-level intervention. This does not replace the legal obligations of regulated entities, but clarifies how operational responsibility should be assigned where different parties contribute to the programmable compliance process.

Governance Functions and Governed Objects

Governance functions may be distributed across multiple parties or combined within a single institution, provided the governed objects remain distinguishable and accountability is not obscured.These functions may be distributed or combined within a single institution, provided they remain distinguishable and accountability over governed objects is not obscured.

Policy artefacts and activation decisions should be attributable, versioned, and supported by a change record sufficient to establish which version was active, where it applied, what dependencies it relied upon, and when it took effect. Authority to activate or update policy should not automatically imply authority to alter trust anchors, eligibility status, or asset-level controls, unless such authority has been expressly assigned.

Table 2: Governance Functions and Governed Objects

Governance function	Governed objects	Typical decisions or actions
Policy governance	Policy artefacts, policy versions, activation states, effective scope, thresholds, escalation conditions	Approves policy logic and parameters; authorises activation, deactivation, and updates; determines effective scope, timing, and conditions of policy changes
Trust dependency governance	Trusted issuers, attestation providers, verification dependencies, trust criteria	Approves, maintains, or removes trusted providers; sets acceptance criteria for attestations and verification sources; reviews ongoing suitability of trust dependencies
Execution governance	Policy wrapper configuration, execution controls, permitted overrides, transaction-handling rules	Operates the compliance-gated execution layer; applies approved policy settings; manages exception handling and override pathways within authorised bounds
Asset governance	Base asset lifecycle controls, embedded administrative capabilities, asset intervention mechanisms	Oversees issuance, redemption, supply-related controls, and any embedded administrative functions; authorises asset-level interventions where applicable
Independent assurance	Audit logs, compliance attestations, decision trails, override records, policy execution outcomes	Reviews governance design and operation; evidences control sufficiency; verifies alignment between policy intent and execution outcomes; supports audit, dispute handling, and supervisory review

Evidence for Material Governance Actions

Material governance actions should generate an evidencing record. These include policy changes such as the activation, deactivation, or modification of policy artefacts or thresholds; eligibility decisions where a participant's ability to hold, receive, or transfer assets is determined; overrides comprising approved exceptions to automated policy evaluation; trust dependency changes including the addition, removal, or modification of trusted issuers, attestation providers, or verification dependencies; and asset interventions such as freezes, administrative transfers, clawbacks, redemptions, or comparable actions.

Each record should identify the action type, the time and relevant transaction or event identifier, the affected asset, wrapper, policy, or trust scope, the actor or authorised source responsible, and the resulting outcome or state change. Additional fields may be captured where needed to explain the basis of the action, such as the policy version applied, the attestation relied upon, or the approving authority.

If implementers choose to anchor evidence references on-chain, such anchoring should be append-only and should not introduce additional execution privileges or blur the distinction between assurance and administration.

Compliance Attestations are execution-time records of compliance evaluation outcomes. They are not equivalent to regulatory filing obligations such as suspicious activity reports, transaction reports, or other regulatory returns, which remain distinct obligations governed by applicable law and reporting requirements[31]. Institutions should ensure that the architecture’s evidencing outputs are clearly distinguished from, and do not purport to substitute for, those statutory reporting obligations.

[31] More broadly, the model does not address, and should not be read as resolving, the question of whether programmable compliance outputs have legal standing equivalent to institutional compliance processes under any jurisdiction's law. Legal enforceability is outside the scope of this paper; that determination rests with regulated institutions, legal counsel, and the applicable regulatory authority.

Conclusion

This paper sets out an architectural model for programmable compliance in tokenised financial systems. It establishes a structured approach to translating regulatory and internal control requirements into governed artefacts, enabling consistent evaluation and enforcement of compliance at the point of execution.

The approach defines how compliance-relevant inputs are assembled, enriched with identity and verification data, evaluated across multiple domains, and applied during execution. This supports consistent enforcement across jurisdictions, asset classes, and operating environments, while allowing policy to evolve independently of core asset logic. It addresses a fundamental tension between frequently changing regulatory requirements and the stability required of execution and asset layers.

Programmable compliance concentrates compliance decision-making at runtime, increasing the consequences of design and governance choices. Deterministic execution and dependence on external data sources are structural features of this environment. The governance model presented here anchors accountability across policy, trust dependencies, execution, asset control, and independent assurance, with sufficient evidence required throughout to support audit, dispute resolution, and supervisory review.

Two areas require further development. The first is technical: common data standards, assessment of behaviour across heterogeneous ledger environments, and mapping to specific deployment contexts where governance arrangements and regulatory obligations differ. The second is institutional and regulatory: deployments spanning jurisdictions will encounter requirements that overlap or diverge, and resolving these requires engagement between institutions, infrastructure operators, and policymakers on questions of policy precedence and how evidencing outputs relate to statutory reporting obligations.

The GL1 initiative envisions an ecosystem of institutional-grade market infrastructures, with programmable compliance as a core component. Practical deployment will inform ongoing refinement across institutions, transaction flows, and institutional roles. The goal is programmable compliance that works consistently across institutions, jurisdictions, and technology environments, and that regulators, institutions, and infrastructure operators can rely on with confidence.

References

1. Bank for International Settlements Innovation Hub, *Project Mandala: Streamlining Cross-Border Transaction Compliance* (Final Report, October 2024), available at: https://www.bis.org/publ/othp87.htm
2. Chainlink, *Automated Compliance Engine (ACE): Technical Overview* (2025), available at: https://blog.chain.link/automated-compliance-engine-technical-overview/
3. European Parliament and Council of the European Union, *Regulation (EU) 2023/1113 of 31 May 2023 on information accompanying transfers of funds and certain crypto-assets and amending Directive (EU) 2015/849* [Transfer of Funds Regulation] (2023), available at: https://eur-lex.europa.eu/eli/reg/2023/1113/oj/eng
4. Fenergo, *Global Financial Institution AML and Regulatory Fines Soar in 2023* (2024), available at: https://resources.fenergo.com/newsroom/global-financial-institution-aml-and-regulatory-fines-soar-in-2023
5. Financial Action Task Force, International Standards on Combating Money Laundering and the Financing of Terrorism & Proliferation: The FATF Recommendations (February 2012, as amended October 2025), available at: https://www.fatf-gafi.org/content/dam/fatf-gafi/recommendations/FATF%20Recommendations%202012.pdf
6. Financial Action Task Force, *Updated Guidance for a Risk-Based Approach to Virtual Assets and Virtual Asset Service Providers* (2021), available at: https://www.fatf-gafi.org/content/dam/fatf/documents/recommendations/Updated-Guidance-VA-VASP.pdf
7. Financial Action Task Force, *Asset Recovery: Guidance and Best Practices* (2025), available at: http://www.fatf-gafi.org/content/dam/fatf-gafi/reports/Asset-Recovery-Guidance-Best-Practices.pdf
8. Financial Action Task Force, *Targeted Update on Implementation of the FATF Standards on Virtual Assets and Virtual Asset Service Providers* (2024), available at: https://www.fatf-gafi.org/content/dam/fatf-gafi/recommendations/2024-Targeted-Update-VA-VASP.pdf
9. Financial Stability Board, *The Financial Stability Implications of Tokenisation* (2024), available at: https://www.fsb.org/uploads/P221024-2.pdf
10. Global Blockchain Business Council and Oliver Wyman, *Proposed Risk Mitigation Framework for Non-Financial Risks of Blockchain Infrastructures:*

Phase 2 Infrastructure (L1/L2) (2026), available at: assets.ctfassets.net/so75yocayyva/66juWKzjrjQoWgV6a7gfo3/5d960a9ce90c393c0ef20b95aa00c4e2/Phase_2_Proposed_Risk_Mitigation_Framework_for_Non-Financial_Risks_of_Blockchain_Infrastructure_2-2.pdf

11. Global Finance & Technology Network, *Shaping the Future of Cross-Border Transaction Compliance* (2023), available at: https://gftn.co/insights/shaping-the-future-of-cross-border-transaction-compliance

12. Global Legal Entity Identifier Foundation, *Introduction to the Verifiable LEI (vLEI): Digital ID for Organisations Everywhere* (Version 2.3, November 2024), available at: https://www.gleif.org/organizational-identity/lei-vlei/the-verifiable-lei-vlei/2024-11-18-gleif_introduction-to-the-vlei_v2.3_compressed.pdf

13. Hong Kong Monetary Authority, *Guideline on Anti-Money Laundering and Counter-Financing of Terrorism for Licensed Stablecoin Issuers* (2025), available at: https://www.hkma.gov.hk/media/eng/doc/key-functions/banking-stability/aml-cft/Guideline_on_Anti-Money_Laundering_and_Counter-Financing_of_Terrorism_For_Licensed_Stablecoin_Issuers_eng.pdf

14. IDEMIA, *Modernizing Digital Identity for Stablecoin Regulation under the GENIUS Act* (2025), available at: https://na.idemia.com/wp-content/uploads/2025/10/Modernizing-Digital-Identity-for-Stablecoin-Regulation-under-the-GENIUS-Act-IDEMIA-white-paper.pdf

15. Institute of Internal Auditors, *The IIA's Three Lines Model: An Update of the Three Lines of Defense* (2020), available at: https://www.theiia.org/globalassets/documents/resources/the-iias-three-lines-model-an-update-of-the-three-lines-of-defense-july-2020/three-lines-model-updated-english.pdf

16. International Organization of Securities Commissions, *Methodology for Assessing Implementation of the IOSCO Objectives and Principles of Securities Regulation* (2011), available at: https://www.iosco.org/library/pubdocs/pdf/ioscopd359.pdf

17. Monetary Authority of Singapore, *Purpose Bound Money (PBM) Technical Whitepaper* (2023a), available at: https://www.mas.gov.sg/-/media/mas-media-library/development/fintech/pbm/pbm-technical-whitepaper.pdf

18. Monetary Authority of Singapore, *Project Guardian Fixed Income Framework* (2023b), available at: https://www.mas.gov.sg/-/media/mas-media-library/development/fintech/guardian/guardian-fixed-income-framework.pdf

19. Monetary Authority of Singapore, *Consultation Paper on Proposed Guidelines on Third-Party Risk Management* (2026), available at:

https://www.mas.gov.sg/publications/consultations/1/consultation-paper-on-proposed-guidelines-on-third-party-risk-management

20. Nasdaq Verafin, *The Impacts of Financial Crime on the U.S. Economy* (2024).

21. Tokeny Solutions, *ERC-3643 Whitepaper: T-REX Protocol* (Version 4, May 2023), available at: https://tokeny.com/wp-content/uploads/2023/05/ERC3643-Whitepaper-T-REX-v4.pdf

22. Wee Kee Toh, Michael Maurer, Emma Landriault, Ashwanth Samuel, Lillian Wang, and Neha Narula, *Designing Payment Tokens for Safety, Integrity, Interoperability, and Usability*, Massachusetts Institute of Technology Digital Currency Initiative and Kinexys by J.P. Morgan, 2025, available at: https://www.dci.mit.edu/projects/designing-payment-tokens-for-safety-integrity-interoperability-and-usability

Acknowledgements

This paper is developed & published by the Global Layer One initiative with contributions from:

- Khai Uy Pham, *Banque de France*
- Sonja Davidovic, *International Monetary Fund*
- Weekee Toh, *Kinexys by J.P. Morgan*
- Kenneth See, *Monetary Authority of Singapore*
- Cayden Chang, *Standard Chartered Bank*
- Sam Vicary, *Standard Chartered Bank*

With special thanks to:

- Jan Philipp Fritsche, *Bermuda*
- Raunak Mittal, *BIS Innovation Hub*
- Edmund To, *Chainlink Labs*
- Ivan Mortimer-Schutts, *GLEIF*