Euro-Native Tokenised Market Infrastructure

Monetary Primacy, Ontology-Driven Control, and Strategic Autonomy

A Response to the Eurosystem's Appia Public Consultation

Zakaryae Boudi
Intelligence Economy Institute
April 2026

This document presents the Intelligence Economy Institute's feedback on the Eurosystem's Appia roadmap, submitted to the European Central Bank's public consultation on tokenised financial market infrastructure. The responses address the high-level principles governing Appia, the design of a shared network layer, European strategic autonomy, common standards, the settlement of central bank money across one or several networks, the governance of Eurosystem core services, and the conduct of the Appia investigation. The strategic framing draws on TKNFRA's work on ontology-driven control and the INTERLOCC settlement protocol, and on a broader analysis of euro monetary sovereignty, digital monetary substitution and euro-native liquidity infrastructure.

1 High-level principles governing Appia (Chapter 2.3)

We support the high-level direction of Appia and would recommend that the Eurosystem make five principles more explicit.

First, monetary primacy should be reflected at architecture level, not only at policy level. In a tokenised environment, the decisive question is not merely whether central bank money remains legally relevant, but whether it remains operationally central in settlement, collateral and liquidity routing. This matters because the emergence of a parallel digital monetary infrastructure, largely dominated by dollar-linked stablecoins, creates a real risk of functional substitution and progressive weakening of euro relevance in transactional and collateral flows.

Second, neutrality should be understood as neutrality of infrastructure rules, not neutrality of sovereignty outcomes. A European tokenised ecosystem should remain open to market participation and innovation, but it should not be indifferent to excessive dependence on non-European infrastructures, non-European governance centres, or non-European liquidity anchors. Strategic autonomy must be embedded as a design objective.

Third, interoperability must be based on semantic discipline, not only on messaging compatibility. It will not be sufficient to connect ledgers at transport level if assets, obligations, compliance states, settlement finality conditions and operational events are not represented through a shared machine-readable model. This is where ontology-driven control becomes critical: it allows market infrastructures, banks and public authorities to reason over common definitions rather than over heterogeneous local abstractions.

Fourth, programmability should be bounded by verifiability. In the next phase of financial market infrastructures, the challenge will not be the existence of programmable logic, but the ability to prove that this logic behaves within defined legal, prudential and operational boundaries. Appia

should therefore elevate formal verification, executable rule models and auditable state transitions from optional technical enhancements to first-order infrastructure principles.

Fifth, resilience should be conceived as layered resilience. This means legal resilience, governance resilience, operational resilience, cryptographic resilience and liquidity resilience. A tokenised market fails not only when systems go offline, but also when participants cannot rely on predictable governance, portable liquidity, clear state models or trusted recovery procedures.

2 The network layer, strategic autonomy and common standards (Chapter 3.2)

2.1 The network layer as a utility

Yes, we believe the concept of a network layer as a utility can provide a strong foundation for a future tokenised financial ecosystem, provided the utility layer is defined carefully. It should not be understood as a monolithic ledger that absorbs every market function. Rather, it should be understood as a neutral coordination substrate for identity, attestations, messaging, settlement orchestration, lifecycle state synchronisation and interoperability.

In our view, the elements most likely to function as a neutral shared utility are common identity and participant qualification layers, shared attestation registries, standardised event and lifecycle semantics, interoperability and routing services, synchronisation of delivery-versus-payment and payment-versus-payment workflows, and audit-grade timestamping and evidence layers. These are domains where public trust, common standards and predictability are more valuable than proprietary fragmentation.

By contrast, product structuring, client servicing, liquidity intermediation, risk transformation, market-making and many forms of balance-sheet deployment should remain competitive layers operated by market participants. The public utility should enable markets, not replace them.

The concept nevertheless requires refinement in three respects. First, the boundary between orchestration and execution must be clarified. Second, semantic interoperability should be explicitly recognised as part of the utility layer. Third, the utility should be designed as a modular stack, so that different market infrastructures and supervised institutions can connect through shared rules without being forced into uniform business logic.

From TKNFRA's perspective, ontology-driven control is particularly relevant here. A utility layer becomes meaningfully neutral only if participants can map different systems into a shared formal representation of assets, obligations, permissions, lifecycle states and controls. Without that semantic layer, "interoperability" risks becoming a narrow transport concept rather than a true market utility.

2.2 Europe's strategic autonomy

Europe's strategic autonomy in a tokenised financial ecosystem will depend on whether it controls not only policy and regulation, but also reference infrastructure, governance mechanisms, semantic standards, cryptographic trust anchors and liquidity corridors. This is especially important in a context where non-European digital dollar infrastructures already benefit from scale and network effects.

Legal autonomy requires that critical settlement logic, rulebooks, compliance enforcement points and dispute-resolution mechanisms remain anchored in European legal frameworks. Technological autonomy requires European capacity in ledger operation, cryptographic infrastructure, secure wallet orchestration, formal assurance and lifecycle control. Operational autonomy requires that Europe can continue to operate and evolve its tokenised infrastructure without dependence on extra-jurisdictional actors that may become vectors of sanction, censorship or unilateral technical leverage. Governance autonomy requires that key rule changes, upgrade processes and participation criteria cannot be captured by commercial coalitions or foreign operational centres.

We would recommend six concrete measures.

The first is to define a European minimum stack for critical tokenised infrastructures, including identity, attestations, settlement semantics, governance procedures, recovery mechanisms and security assurance requirements.

The second is to require semantic and legal portability. Participants should be able to migrate positions, collateral references and settlement relationships across compliant networks without losing legal clarity or operational continuity.

The third is to establish sovereign governance requirements for core shared components, including transparent change-management and public-interest veto rights where central bank services are implicated.

The fourth is to make formal assurance part of autonomy. An ecosystem that depends on opaque software supplied by a small number of actors is not strategically autonomous even if it is hosted in Europe. Formal models, verifiable modules and auditable lifecycle transitions reduce hidden dependence.

The fifth is to develop euro-native liquidity infrastructures rather than relying on foreign-denominated liquidity pools. Autonomy is not only defensive regulation; it is also proactive liquidity architecture.

The sixth is to ensure that strategic interoperability with external ecosystems does not imply internal dependence on them. Europe should interoperate outward from a position of architectural coherence, not enter tokenised finance as a downstream user of external standards and rails.

2.3 Common standards, rules and practices

The best way to establish common standards, rules and practices is to combine public coordination with modular technical implementation. Europe should avoid both extremes: excessive abstraction without implementation discipline, and fragmented vendor-led standards that never converge at ecosystem level.

Standardisation should be pursued in at least six domains.

The first domain is asset and cash semantics: what an asset is, who may hold it, how lifecycle events are represented, and how legal claims are mapped to digital state transitions.

The second is compliance and attestations: how KYC, AML, sanctions screening, eligibility, holding constraints and jurisdictional permissions are represented and verified.

The third is settlement semantics: finality, conditionality, DvP, PvP, netting, reservation, release, cancellation, exception handling and evidence.

The fourth is identity, role and authority models: institutions, custodians, operators, validators, overseers and service providers should interact under standardised role definitions.

The fifth is governance: upgrade rights, emergency interventions, policy change processes, auditability and dispute management.

The sixth is observability and reporting: event logs, supervisory views, machine-readable reporting and evidence formats.

Public institutions should play three roles. They should act as conveners of minimum common semantics and policy requirements. They should act as certifiers of baseline trust conditions for infrastructures that want access to public money interfaces. And they should act as stewards of the common layer where fragmentation would create systemic inefficiency or prudential ambiguity.

In our view, standards should be machine-interpretable from the start. This is why ontology-driven control is useful: it allows standards to exist not only as documentation, but as operational artefacts that govern validation, interoperability and assurance across infrastructures.

3 Settlement of central bank money and core services (Chapter 3.3)

3.1 Settlement on a single network

Enabling central bank money settlement on one network has the advantage of maximal coherence. It can simplify finality logic, reduce ambiguity, concentrate liquidity, improve supervisory visibility and create a powerful focal point for standardisation and market coordination. It may also accelerate the emergence of a strong euro-native reference environment.

However, the trade-offs are material. A single network model creates concentration risk in governance, technology, operations and innovation pathways. It may discourage market experimentation, slow ecosystem adaptation and create a perception that access to central bank settlement is tied to one technical architecture. Even if that network is resilient, the strategic dependence created by singularity may itself become a weakness.

There is also a political economy issue. In a rapidly evolving tokenised ecosystem, no single technical environment is likely to remain optimal across all use cases, asset classes and market structures. A one-network approach may produce near-term clarity but longer-term rigidity.

Our view is therefore nuanced. One network can be appropriate as a lighthouse model or primary reference environment, especially in early phases. But if adopted, it should be architected as a modular and semantically open environment, with clear pathways for future interconnection and portability. In other words, one network should not mean one closed universe.

3.2 Settlement on a limited number of selected networks

Settlement on a limited number of selected networks appears to us the more balanced long-term approach. It offers resilience through controlled diversity, reduces single-point concentration risk, preserves innovation incentives and makes it possible to align central bank money services with different market structures, provided selection criteria and interoperability obligations are strict.

The principal risk in a multi-network approach is fragmentation. Liquidity may disperse, standards may diverge and supervisory complexity may increase. Yet fragmentation is not an inevitable consequence of plurality. It becomes problematic mainly when networks are not bound by common settlement semantics, governance expectations, portability rules and evidence standards.

For that reason, we would support a limited-network model only if the Eurosystem defines a common semantic and assurance layer across eligible networks. This is precisely where a formal, ontology-based approach can be decisive. If all selected networks must expose compliant representations of cash states, asset states, permissions, attestations and settlement events, then network plurality need not mean legal or operational incoherence.

This model would also be more compatible with the industrial objective of preserving European innovation capacity while avoiding uncontrolled infrastructural drift.

3.3 Core services on a directly operated network

There is a strong case for the Eurosystem providing core services on a network it operates directly, especially where public trust, settlement integrity, monetary primacy and baseline resilience are concerned. A directly operated network could serve as the sovereign anchor of the ecosystem, a benchmark for governance and a reference point for interoperability.

To support business needs efficiently, such a network would need several features. It would require deterministic finality semantics, auditable lifecycle events, strong identity and permissioning models, robust privacy controls for institutional workflows, interoperable APIs and messaging interfaces, programmable but bounded settlement conditions, transparent rule-change procedures, and explicit support for tokenised cash interaction with tokenised securities and other financial claims.

It would also need to support sophisticated market workflows rather than only simple trans-

fers. This includes conditional settlement, reservation and release logic, atomicity patterns, netting-aware flows, attestation-based progression of state and structured exception handling. Our INTERLOCC work is relevant here because it is precisely focused on configurable institutional settlement orchestration, including direct on-chain settlement, deferred settlement and netting-based configurations under auditable lifecycle control.

Finally, if the Eurosystem operates the network directly, it should still allow verified external modules, standardised connectors and semantic interoperability. Direct operation should guarantee trust and coherence, not exclude innovation.

3.4 Core services under shared governance

Shared governance can be beneficial if Europe wants to combine public trust with market dynamism and technical pluralism. It can mobilise expertise across central banks, market infrastructures, supervised institutions and technical operators. It can also improve adoption by giving participants a meaningful role in ecosystem evolution.

But shared governance only works if it is deeply structured. In infrastructures involving central bank money, informal consortia or vendor-led governance are not sufficient. The governance model must define decision rights, veto thresholds, emergency procedures, transparency requirements, liability allocation, audit rights, upgrade processes, admission criteria and conflict-resolution paths with precision.

From our standpoint, three conditions are essential. First, public-interest safeguards must remain dominant where settlement integrity and monetary functions are concerned. Second, governance must be machine-enforceable wherever possible, not merely descriptive. Third, all change processes should be evidence-rich and reviewable, so that no participant is exposed to opaque infrastructure evolution.

This is again an area where ontology-driven control and formal specification can add practical value. Governance becomes safer when roles, rights, conditions and allowed state transitions are encoded in a structured and reviewable way, rather than dispersed across prose, code and committee interpretation.

4 The Appia investigation (Chapter 4)

4.1 General approach to the investigation (Chapter 4.1)

We would encourage the Eurosystem to structure the Appia investigation around executable use cases rather than around abstract architecture alone. The most useful path would be to test the ecosystem through concrete workflows where tokenised cash, tokenised assets, attestations, compliance events and settlement conditions interact under realistic constraints.

In particular, the investigation should include cross-border payments, tokenised securities settlement, collateral mobility, conditional funding, institutional treasury management and netting-enabled interbank flows. These are the use cases where the euro's role as transactional, collateral and coordination asset will actually be tested.

We would also recommend that Appia distinguish clearly between infrastructure validity and economic viability. A design may be technically sound and still fail if it does not attract liquidity, reduce friction or support credible market incentives. Infrastructure must be paired with liquidity strategy, otherwise network effects will reinforce already dominant non-euro rails.

Lastly, the investigation should include assurance by design. Formal models, ontology-backed semantic specifications and verifiable state machines should be included early, because they are much harder to retrofit once architectures and market practices have already fragmented.

4.2 Public-private collaboration (Chapter 4.2)

Efficient public-private collaboration under Appia will require more than consultation. It will require controlled co-production. Public authorities should define the policy perimeter, public-interest constraints and core trust requirements. Private actors should help test implementation pathways, operational models and market incentive structures. But the interaction should be organised around common artefacts: shared semantic models, test scenarios, assurance criteria, interoperability profiles and evidence frameworks.

We would recommend three forms of collaboration.

The first is supervised design labs focused on specific use cases, with public and private participants working from a common specification base.

The second is reference implementations and controlled pilots that test both technical performance and governance assumptions.

The third is an assurance forum where formal methods, auditability, semantic consistency and operational evidence are treated as common concerns, not left to each participant individually.

This is important because tokenised finance is not only an engineering problem. It is also a coordination problem between law, market practice, software logic and public trust. Collaboration should therefore be organised in a way that produces reusable European public goods.

5 Additional feedback

The central strategic issue behind Appia is not simply whether Europe can participate in tokenised finance, but whether it can do so from a position of monetary and infrastructural relevance. The risk is not abstract. If digital financial activity progressively settles, collateralises and scales on non-euro rails, Europe may preserve regulatory authority in principle while losing practical influence over the architecture of digital finance.

Appia should therefore be treated not only as a technical initiative, but as an industrial and monetary coordination project. Europe needs sovereign market infrastructure, euro-native liquidity pathways, machine-readable standards, and assurance frameworks that make public trust compatible with programmable finance.