

September 17, 2026

Bank of England
Financial Market Infrastructure Directorate
Threadneedle Street
London EC2R 8AH

Re: Consultation on the Draft Code of Practice for Sterling-Denominated Systemic Stablecoin Issuers

Response deadline: 22 September 2026

SYSTEMIC STABLECOINS AND THE INFRASTRUCTURE ON WHICH THEY DEPEND

Operational Evidence, Settlement Integrity and Financial Stability

To the Bank of England:

TauriPay, a Public Benefit Corporation respectfully submits this response to the Bank of England's consultation on the draft Code of Practice for sterling-denominated systemic stablecoin issuers.¹

TauriPay supports the Bank's objective of allowing new forms of digital money to develop while preserving confidence in money and UK financial stability. This response focuses on the infrastructure evidence required when fulfilment of a systemic stablecoin issuer's obligations depends materially upon software-defined, distributed, layered or otherwise external transaction infrastructure.

Backing is necessary, but it is not sufficient to establish suitability as systemic financial infrastructure.

The central question addressed by this response is not whether a particular architecture should be permitted, but what should be demonstrated about infrastructure before a stablecoin materially dependent upon it is relied upon for a systemic monetary function.

¹ Bank of England, Sterling-denominated systemic stablecoins: Policy Statement and consultation on draft Code of Practice (22 June 2026).

Executive Summary

The Bank's proposed regime establishes important requirements concerning backing assets, safeguarding, capital and reserves, legal claim, redemption and the temporary issuance guardrail, while indicating that the broader supervisory framework will continue to develop through operational-resilience, reporting, record-keeping and FMI work.²

This response starts from existing prudential principles rather than proposing a substitute framework. CPMI-IOSCO has confirmed that systemically important stablecoin arrangements used for payments are expected to observe the relevant Principles for Financial Market Infrastructures, including attention to governance, comprehensive risk management, settlement finality and money settlements. The Financial Stability Board likewise calls for comprehensive oversight of global stablecoin functions, clear governance and accountability, effective risk management and timely redemption.³⁴

The additional question is how those established outcomes can be demonstrated when performance of the regulated function depends upon transaction infrastructure with operating characteristics distinct from the issuer's reserve portfolio or legal structure.

Before infrastructure becomes a material dependency of systemic money, the evidence supporting its fitness for purpose should be commensurate with the financial function, scale and stress conditions it may be required to support. That assessment should follow the complete dependency and assurance path, including the boundaries at which authority, observation, responsibility or assurance can change. As the infrastructure, workload or dependency chain changes, the conclusion should remain subject to continuing evidence and reassessment.

The response makes six propositions, all directed to one supervisory question: what evidence is sufficient to justify reliance on infrastructure supporting a systemic monetary function?

The analysis can be reduced to four questions: What financial state exists? Who or what can establish, observe or change it? What evidence makes that state safe to rely upon? What happens when the systems carrying the obligation disagree, degrade or fail? The technical sections that follow are intended to make those questions operational for supervision, not to create a separate technology framework.

Appendix H applies those questions through an illustrative, proportionate set of assessment questions for arrangements that materially depend on distributed infrastructure. It is intended to help identify where state is established, how it is observed or represented, what evidence supports reliance, and how the relevant financial state can be reconstructed and reconciled following disruption.

For supervisory purposes, the analysis can be organised around five capabilities that should mature as systemic reliance increases:

Supervisory capability	Question to be answered	Principal sections
1. Dependency and boundary mapping	Which systems, services and institutional actors are material to the financial function, and what proposition crosses each boundary?	II, III, VIII
2. State and authority determination	Which source is authoritative for the relevant proposition, what is observation only, and what remains unresolved?	VI, VII, VIII
3. Evidence and production fitness	What evidence justifies reliance at the contemplated scale, assurance threshold and stress condition?	IV, V, XI
4. Stress, recovery and state reconstruction	Can the arrangement distinguish failure from indeterminate outcome, prevent unsafe progression, and reconstruct authoritative financial state?	VIII, IX
5. Continuing supervisory assurance	What evidence must remain available over time, and what material change should trigger reassessment?	X, XI, XII, XIII

² Bank of England, Sterling-denominated systemic stablecoins: Policy Statement and consultation on draft Code of Practice (22 June 2026).

³ Financial Stability Board, High-level Recommendations for the Regulation, Supervision and Oversight of Global Stablecoin Arrangements (17 July 2023).

⁴ CPMI-IOSCO, Application of the Principles for Financial Market Infrastructures to stablecoin arrangements (13 July 2022).

1. Systemic stablecoins should be evaluated across the complete operational lifecycle. A stablecoin transfer or redemption may involve authorisation, signing, submission, ordering, execution, confirmation, finality, issuer recognition, reconciliation and fiat settlement. Each material stage may have its own dependency, timestamp, control point and failure condition.
2. Instantiation creates relevant operational dependency. Where issuance, transfer, settlement, redemption, ownership determination or authoritative state materially depends upon particular infrastructure, the characteristics of that infrastructure that can affect those functions become relevant to the end-to-end risk profile of the arrangement.
3. **Technology neutrality does not mean infrastructure blindness. Public, private, permissioned, permissionless, layered and conventional architectures should be evaluated according to the outcomes they demonstrably produce and the dependencies upon which those outcomes rely.**
4. **Functional demonstration is not production fitness.** A successful pilot can prove that a workflow functions; a component benchmark can prove the performance of a particular component. Neither necessarily proves full-lifecycle capacity or resilience under representative production, stress, fault or, where material, adversarial conditions.
Capacity is a necessary condition for systemic reliance: where infrastructure is shared, sufficient sustainable successful capacity should remain available to the regulated financial function at the required assurance level, including when multiple workloads draw on the same capacity during periods of stress.
5. Systemic recognition should expand the evidentiary perimeter. As scale, interconnectedness and criticality increase, the issuer should be capable of demonstrating that material infrastructure dependencies required to discharge regulated obligations satisfy proportionate operational outcomes.
6. **Systemic confidence depends upon the ability to establish material operational facts under stress. Supervisors and institutions may need to know, within a decision-useful timeframe, what succeeded, failed, remains pending, was retried, became final, was changed by privileged action, and which state is authoritative.**

This framing has direct relevance to redemption. The Bank's proposed regime requires direct and timely redemption and links the redemption clock to operational facts, including receipt of the stablecoin by the issuer. If congestion, finality ambiguity, sequencer disruption, bridge failure or another material dependency impairs the holder's ability to deliver the stablecoin into the redemption process, infrastructure behaviour can affect the practical exercise of a prudential right.

The paper therefore distinguishes the financial layer of stablecoin risk from the infrastructure layer beneath it. Reserve sufficiency and liquidity remain essential. But a financially sound issuer can still face operational constraints if the infrastructure required to transfer, recognise, redeem, reconcile or reconstruct the instrument is degraded.

The recommendations are deliberately proportional. A pilot does not require the same evidence as systemic infrastructure. The evidence standard should mature with reliance:

ASSURANCE MATURES WITH RELIANCE

The evidence standard should scale with the criticality of the financial function.




Figure 1. Assurance Matures With Reliance

Implementation principle: the Bank need not certify underlying protocols or prescribe architecture. Where compliance with a prudential outcome materially depends on external infrastructure, the issuer should be able to demonstrate a proportionate evidentiary basis for that reliance, and the Bank should be able to assess the resulting dependency and residual risk.

Accordingly, the response proposes three measures: (1) preserve complete-lifecycle and material-dependency analysis in the current Code; (2) develop proportionate evidence expectations for capacity, finality, failure, recovery, control and measurement provenance as systemic reliance grows; and (3) require reassessment where material change alters a property on which prior suitability depended.

Contents

Executive Summary.....	2
I. The Bank's Framework and the Additional Infrastructure Question.....	5
II. Instantiation and the Infrastructure Dependency of Systemic Money.....	7
III. The Complete Stablecoin Lifecycle and Redemption-Effective Capacity.....	9
IV. Observed Behaviour and Demonstrated Operational Evidence.....	11
V. Capacity, Measurement Provenance and Production Fitness.....	13
VI. Settlement, Finality, Reversibility and Authoritative State.....	16
VII. Governance, Control and Effective Decentralisation.....	18
VIII. Third-Party Dependencies, Operational Resilience, Recovery and Runbooks.....	20
IX. The Systemic Stress Pathway: When Infrastructure Stress Becomes Confidence Stress.....	23
X. Systemic Recognition Should Expand the Assurance Perimeter.....	26
XI. Supervisory Observability and Continuing Operational Evidence.....	27
XII. Proportionality, Innovation and Technology Neutrality.....	29
XIII. Recommendations.....	31
XIV. Conclusion.....	33
Appendix A. Infrastructure Performance Evidence and Illustrative Measurements.....	34
Appendix B. Complete Systemic Stablecoin Transaction and Redemption Lifecycle.....	37
Appendix C. Measurement Provenance Framework.....	38
Appendix D. Fit-for-Purpose Assurance Framework.....	39
Appendix E. Systemic Stress Scenarios and Evidence Requirements.....	41
Appendix F. Consultation Question Crosswalk.....	42
Appendix G. Prudential Standards Crosswalk.....	43
Appendix H. Questions Before Systemic Reliance.....	45

I. The Bank's Framework and the Additional Infrastructure Question

Relevant consultation questions: Q1, Q2 and Q4

I-A. This Response Builds Upon Existing Prudential Principles

This response does not ask the Bank to create a separate technology regime. It asks that, where an established prudential outcome depends materially on external or software-defined infrastructure, the evidence supporting that outcome follow the dependency far enough to show that the complete financial function can actually be relied upon.

The starting point is established prudential policy. CPMI-IOSCO guidance applies the PFMI to systemically important stablecoin arrangements used for payments and clarifies that the arrangement should be considered functionally, including entities integral to it. The FSB's global stablecoin recommendations similarly emphasize comprehensive oversight of functions and activities, clear accountability, effective risk management and timely redemption. The CPMI-IOSCO guidance applies the PFMI to a systemically important stablecoin arrangement performing a transfer function, including entities integral to that arrangement; it does not create a separate stablecoin prudential standard.

The issue addressed here is therefore not whether systemically important stablecoins should satisfy prudential outcomes comparable to other critical financial arrangements. That principle is substantially established. The question is how those outcomes can be demonstrated when the regulated function depends materially upon distributed, layered or otherwise software-defined infrastructure.

This response is methodological. It seeks to connect established prudential outcomes to observable evidence concerning the infrastructure upon which those outcomes rely. The focus is operational behaviour rather than architectural claims, while preserving existing supervisory objectives and improving the organisation and corroboration of evidence.

I-B. Two Analytically Distinct Layers of Stress

A stablecoin arrangement can face issuer and asset-side stress: rapid redemption demand, reserve liquidity, capital adequacy, safeguarding, continuity and wind-down. It can also face transaction-infrastructure stress: capacity saturation, failures, retries, ordering uncertainty, delayed finality, sequencer or validator disruption, bridge or access-layer impairment, privileged intervention, inconsistent records and inability to reconstruct authoritative state.

These layers interact. A financially sound issuer may still have difficulty satisfying redemptions if holders cannot deliver coins reliably into the redemption workflow. Infrastructure disruption can itself increase uncertainty and redemption demand. Elevated redemptions can increase transaction load at precisely the moment the infrastructure is already degraded.

The prudential assurance perimeter should follow the material functions and dependencies required to produce the regulated outcome.

I-C. Scope and Allocation Across the Developing Framework

TauriPay does not recommend placing every technical requirement discussed in this response into the current draft Code. The better approach is to preserve the relevant issuer obligations and prudential outcomes in the Code, while developing detailed operational evidence, reporting, resilience and PFMI expectations through the broader supervisory framework as the Bank's regime matures.⁵

The distinction matters because it avoids both extremes: treating infrastructure as merely an implementation detail, or attempting to regulate every technological component directly. A systemic issuer should remain accountable for identifying and understanding the material dependencies upon which its regulated obligations rely and for producing sufficient evidence concerning those dependencies. Where assurance properly sits at arrangement, recognised-payment-system, service-provider or broader PFMI level, the supervisory framework can allocate the corresponding requirement to the entity or function best placed to satisfy it.

I-D. What This Adds to Existing Operational-Resilience Requirements

The Bank's existing operational-resilience framework already provides an important foundation. For recognised payment system operators and specified service providers, it requires identification of important business services, mapping of the resources and dependencies that support them, setting of impact tolerances, and testing against a range of extreme but plausible disruption scenarios.⁶

This response does not seek to duplicate that framework. The existing framework provides the supervisory structure into which these questions can fit: end-to-end mapping, interdependency analysis, impact tolerances, scenario testing, recovery and governance. What is not yet explicit is the state, authority and assurance methodology needed to apply that structure

⁵ Bank of England, Sterling-denominated systemic stablecoins: Policy Statement and consultation on draft Code of Practice (22 June 2026).

⁶ Bank of England, Operational resilience: recognised payment system operators and specified service providers, Supervisory Statement (12 June 2026), paras 3.14-3.21.

to composite software-defined financial infrastructure. Conventional dependency mapping can establish which components support an important business service and how disruption affects service delivery. Where a financial workflow depends on multiple technical and institutional domains, a further question is what proposition each boundary establishes, which source is authoritative for that proposition, how it is observed, what remains unresolved and what action is safe to take next.

The distinction matters most when components remain individually available but the financial state is partial, differently observed or indeterminate. In that circumstance, conventional service-health indicators can remain within expected ranges while the arrangement has lost a coherent basis for deciding whether a transfer, redemption or settlement action can safely progress. The framework proposed here should therefore be read as a state-and-assurance methodology for applying established end-to-end operational-resilience principles to composite financial infrastructure, not as a parallel resilience regime.

Recent joint FCA/Bank of England work on wholesale-market tokenisation reinforces the relevance of this distinction. The September 2026 Feedback Statement retains technology-agnostic regulatory principles while identifying operational and cyber-resilience concerns in areas including bridges, key-management systems, governance arrangements, oracle providers and cross-chain messaging layers. The point for this response is not that those components are inherently unsuitable, but that architecture-specific dependencies can become material to the regulated outcome and should therefore be visible in the assurance analysis.⁷

⁷ Financial Conduct Authority and Bank of England, FS26/1: Tokenisation in Wholesale Financial Markets (14 September 2026), pp. 9-10.

II. Instantiation and the Infrastructure Dependency of Systemic Money

Supervisory relevance. The relevant question is where the financial state is established and which systems must function correctly for a holder to exercise the monetary claim.

Relevant consultation questions: Q2, Q3 and Q4

II-A. Instantiation Creates Relevant Operational Dependency

A stablecoin does not operate independently of the infrastructure through which it is issued, held, transferred, settled or redeemed. Where issuance, transfer, settlement, redemption, ownership determination or authoritative state materially depends upon particular infrastructure, the characteristics of that infrastructure that can affect those functions become relevant to the end-to-end risk profile of the arrangement.

Where an infrastructure property can materially affect the ability of a systemic stablecoin to be issued, transferred, settled, redeemed, recovered or reconciled, that property is relevant to the prudential analysis. This does not imply that the instrument inherits every technical property or risk of the infrastructure on which it is instantiated.

Instantiation does not mean that technology determines the legal character of the instrument, nor that every technical characteristic automatically becomes a prudential concern. It means that an economically equivalent claim can be exposed to different ordering, capacity, failure, finality, governance and recovery characteristics depending upon the infrastructure required to exercise the claim.

II-B. Instantiation and Authoritative Operational State

The significance of infrastructure dependency can extend beyond service availability. In some arrangements, the distributed infrastructure participates directly in establishing the operative state of the instrument: whether it has been issued, who controls it, whether it has been transferred, and whether it can be delivered into another financial process.

Institutions may maintain internal books, custodial records, indexers, digital twins or other synchronized representations of that state. Those records can be important for reconciliation and recovery, but they do not necessarily provide an independent mechanism for producing a valid subsequent state transition.

A synchronized representation of state is not necessarily an independent mechanism for changing state.

This distinction becomes relevant during disruption. Where the infrastructure required to establish or change the operative state is unavailable, congested or uncertain, parallel institutional records may assist reconstruction without necessarily restoring transferability, settlement or redemption functionality.

Conventional financial infrastructure often places the establishment, observation and control of business state within a bounded, accountable operating domain. In distributed or composite arrangements, those functions can separate: one system may establish state, another may observe or represent it, another may determine whether it is sufficient for a business action, while legal or operational responsibility may remain elsewhere. The prudential issue is therefore not decentralisation itself, but whether those authority, observation and reliance boundaries remain explicit and sufficiently assured as the financial function crosses them.

The prudential question is therefore not whether a distributed ledger is legally designated the single or final record in every architecture. It is whether performance of the regulated financial function materially depends upon state established or changed through that infrastructure. Where it does, the availability, finality, recoverability and authoritative-state characteristics of the infrastructure become relevant to the end-to-end prudential assessment.

Where the exercise, transfer or recognition of a financial claim depends upon distributed infrastructure, the relevant ledger state is not merely an operational record. It can form a material part of the evidence by which the arrangement determines what financial state exists, who is entitled to act upon it and whether the next action can safely proceed. The prudential significance therefore lies not in treating the ledger as the legal instrument itself, but in understanding when the regulated outcome depends upon the integrity, authority, observation and recoverability of the state that the infrastructure establishes or represents.

For distributed financial infrastructure, this means that prudential assurance may ultimately need to reach the state-transition layer itself: the rules and mechanisms by which valid state can change, transactions are admitted and ordered, conflicting outcomes are prevented or resolved, and committed state becomes authoritative. Application logic, custody controls, orchestration and synchronized records can mitigate important risks, but they do not by themselves establish the behaviour of the underlying state-transition system on which the financial function materially depends.

II-C. Architecture Labels Are Not Prudential Conclusions

Labels such as public, private, permissioned, permissionless, Layer 1, Layer 2, distributed or centralised can assist description. They do not establish fitness for purpose.

The proposed framework is therefore complementary to established prudential standards: it does not replace their outcomes, but makes the evidentiary connection between those outcomes and the actual architecture, dependencies and complete financial lifecycle more explicit.

A public network may offer extensive production history and broad participation while presenting particular questions concerning congestion, finality, governance, operator concentration or dependencies. A private system may offer explicit governance and participant controls while presenting different questions concerning concentrated authority, limited production history, vendor dependency or performance evidence derived primarily from controlled environments.

Supervisory relevance. The same legal claim can depend on different mechanisms for creating, changing and recovering its operative state. An issuer-controlled database may concentrate state production, access and recovery within one accountable operator; permissioned, public and layered arrangements can distribute or add ordering, validation, administration, upgrade, proof, data-availability, bridge or settlement dependencies. None is inherently suitable or unsuitable. Each creates a different reliance proposition that should be understood before systemic reliance develops.

Architecture pattern	Where operative state may be established	Questions that become material
Issuer / conventional	Issuer or controlled database and payment infrastructure	Operator control, access, redundancy, change management, recovery, reconciliation
Permissioned / institutional DLT	Shared state among governed participants	Membership, quorum, administrator powers, common infrastructure, upgrades, recovery authority
Public L1	Public transaction and consensus environment plus issuer logic	Ordering, finality, congestion, governance, access dependencies, external operator influence
Layered / composite	Multiple execution, observation, assurance and settlement domains	Sequencer, proof/challenge, data availability, bridge, L1, orchestration, cross-domain state coherence

Figure 2. Where the financial state lives changes the dependency and assurance questions, not the prudential objective.

The appropriate inquiry is functional: which design patterns and operational dependencies does the arrangement actually inherit, and what evidence demonstrates that they are compatible with the contemplated systemic use?

II-D. Material Dependency Defines the Assurance Perimeter

For purposes of this response, a material infrastructure dependency is a system, component, service or control function whose failure, degradation, delay, intervention or material change could materially impair issuance, access, ownership determination, transfer, settlement, redemption, reconciliation, recovery or authoritative state.

Materiality avoids indiscriminate regulation. A peripheral analytics service may be immaterial. A sequencer, bridge, custody system, key-management process, transaction-access layer or underlying ledger that can prevent holders from transferring coins into the redemption process may be highly material.⁸

The map should evolve as the architecture evolves. Systemic recognition should therefore be accompanied by explicit identification of the material dependencies upon which fulfilment of the issuer's regulated obligations relies. Appendix H provides illustrative questions for applying this inquiry where distributed state, layered execution or dependent representations are material.

INSTANTIATION AND AUTHORITATIVE OPERATIONAL STATE

Why infrastructure characteristics can become part of the end-to-end prudential assessment.




Figure 3. Instantiation and Authoritative Operational State

⁸ Financial Conduct Authority and Bank of England, FS26/1: Tokenisation in Wholesale Financial Markets (14 September 2026), p. 10.

III. The Complete Stablecoin Lifecycle and Redemption-Effective Capacity

Supervisory relevance. A holder experiences one monetary claim, but performance of that claim can depend on multiple technical and institutional steps. Evidence confined to one stage can therefore miss the point at which the financial obligation is delayed, duplicated, left unresolved or prevented from completing.

Relevant consultation questions: Q1, Q2, Q3 and Q4

III-A. A Stablecoin Transaction Is Not One Event

Stablecoin activity is often described through intuitive states such as sent, confirmed, settled or redeemed. For supervisory and operational purposes, those labels can collapse multiple events.

**Intent → Approval → Signed → Submitted → Ordered → Included → Executed → Confirmed → Finalised
→ Issuer Recognised → Bank Settled**

A more complete lifecycle can also include custody and key-policy approval, access-layer receipt, gateways or adapters, RPC or node access, orchestration and routing, sequencer or mempool receipt, ordering, application or smart-contract execution, data-availability or proof dependencies, state transition, bridge or cross-domain processing, issuer or custodian book updates, redemption orchestration, bank or payment-system settlement, reserve-side settlement and post-event reconciliation. The number and sequence of material stages will vary by architecture, which is why the lifecycle should not be collapsed into a single word such as 'instant.'

Descriptions such as direct, instant or atomic can therefore simplify the visible transaction without necessarily eliminating the functions beneath it. Depending on the architecture, authority may instead be redistributed across custody and signing, access, sequencing, validation, cross-domain attestation, observation, issuer recognition, reconciliation and cash settlement. The relevant supervisory question is whether those boundaries have actually been removed or merely relocated behind a simplified interface.

No single component necessarily observes, controls or establishes the complete financial proposition. A custody system may establish authorisation without establishing execution; a gateway may establish receipt and routing without establishing destination outcome; an observer may report a state without being authoritative for it; an orchestrator may determine workflow action while relying on another system for the underlying financial fact; and a bridge or cross-domain mechanism may establish one domain's state while another dependency remains outstanding.

The prudential issue is therefore not only whether each component is available. It is whether the propositions passed across component boundaries remain sufficiently authoritative, observable, correlated and assured for the next financial action. A composite arrangement can contain individually functioning components and still lose coherence around the state of the financial obligation.

The complete lifecycle is also the relevant structural-resilience perimeter. A materially adverse outcome need not require compromise of the base ledger or consensus mechanism itself. It can arise through compromise, manipulation or inconsistent behaviour at a custody or key boundary, transaction-access point, ordering or sequencing function, validation process, bridge or cross-domain mechanism, observation path, orchestration layer, privileged control or recovery process. The prudential question is whether the financial propositions required by the workflow remain preserved, observable and safely bounded when one or more material boundaries behaves incorrectly or adversarially.

Each material stage can have its own timestamp, dependency, control point, failure mode and legal or economic consequence. For a systemic arrangement, supervisors and institutions should be able to determine when an instruction was authorised, accepted, ordered and executed; what state resulted; when that state became final; when the issuer recognised receipt; when redemption became valid; and which record is authoritative where systems disagree.

Progress through a technical lifecycle does not necessarily establish that responsibility has transferred or that the resulting state has reached the assurance threshold required for the financial obligation. State advance, responsibility transfer and assurance can coincide in some arrangements and diverge in others; where they diverge, the distinction should remain visible to the issuer, counterparties and supervisor.

One way of analysing this timing relationship is to distinguish a business clock, an infrastructure-processing clock and an assurance/finality clock. A systemic stablecoin arrangement may also operate across several non-coincident clocks: a business clock governing the financial obligation, an infrastructure-processing clock governing submission, ordering and execution, and an assurance clock governing when the resulting state becomes sufficiently established for the relevant reliance. The prudential issue is whether those clocks can diverge without making the obligation indeterminate, unperformable or unreconstructable.

A related but distinct requirement is certainty of outcome (or certainty of fate): whether a participant can determine reliably that an instruction succeeded, failed, or remains unresolved. A determinable failure can enter a controlled exception, reconciliation or retry process; an indeterminate outcome may require authoritative-state reconciliation before another economic action is taken, because a business retry, native resubmission or replacement attempt can create different technical effects against the same underlying financial instruction. Certainty of outcome need not coincide with technical or

legal finality, but where the states differ the distinction should be explicit because unresolved outcomes can drive reattempts, liquidity uncertainty, duplicate instructions and inconsistent operational treatment.

III-B. Redemption Creates an End-to-End Infrastructure Requirement

The Bank's redemption requirements make this lifecycle operationally relevant. The issuer may owe a legal obligation while relying on external infrastructure for the holder to deliver the instrument into the redemption process.⁹

Under stress, redemption demand and transaction demand may rise together. Latency can increase; transactions can fail or remain unresolved; wallets and intermediaries may resubmit instructions; offered load can grow faster than economic demand; and the issuer may have to determine which transfers are complete, final and eligible for recognition.

The reserve portfolio can remain sound throughout this sequence. The legal claim can remain valid. Yet infrastructure behaviour can determine whether the holder can practically exercise the redemption right.

This is why temporary congestion can become more than an operational inconvenience. The infrastructure does not change the existence of the legal redemption claim. But where a holder cannot reliably deliver the instrument, establish the relevant state or have receipt recognised by the issuer within the applicable period, a legally valid right may not translate into timely economic access to the underlying money. The prudential issue is therefore the practical performance of the claim, not its legal existence.

This point does not seek to redefine the Bank's proposed 24-hour redemption timeframe. Under the June 2026 proposal, the clock begins once the issuer has received the redemption request, completed the relevant AML/KYC checks and received the stablecoins in its wallet, and it stops once the issuer or a third party acting on its behalf has instructed a valid payment order. The separate operational-access question is whether infrastructure required to deliver the coins into that process can itself become a material constraint on continuous redeemability, particularly under stress.¹⁰

III-C. Redemption-Effective Capacity

Traditional capacity measures ask how many transactions the infrastructure can process. For systemic stablecoins, that is incomplete.

Redemption-effective capacity is the sustainable rate at which the complete arrangement can carry valid redemption activity from holder initiation through the material infrastructure stages required for issuer receipt and recognition while preserving required finality, reconciliation and operational-control outcomes.

This is not intended to replace conventional throughput metrics. It connects infrastructure performance to the regulated function.

Useful evidence may need to distinguish among gross instructions submitted, instructions accepted, successful execution, failed or unresolved transactions, retries, successful finalisation, delivery into issuer-controlled infrastructure, issuer-recognised redemption requests and redemptions successfully completed.

For systemic use, relevant evidence can therefore include successful finalised throughput, failure and retry rates, retry-adjusted throughput and finality-to-redemption-recognition latency.

For prudential purposes, the more useful capacity concept is capacity at the required assurance level. Claimed or benchmark throughput, observed gross activity, successful activity, assured or finalised activity and complete regulated outcomes are different propositions. Redemption-effective capacity is the function-specific expression of this principle for the Bank's regime: how much valid redemption workload can reach the state on which the financial obligation relies while required assurance, control and recovery properties remain within acceptable bounds.

III-D. Capacity Should Be Understood Before Critical Reliance Develops

Systemic capacity should not first be discovered during a run. A benchmark tells you what a system did. Capacity planning asks what it must be able to do. For systemic use, that planning should reflect the required operating envelope, including representative peaks, stressed conditions, degradation, retries and recovery while preserving the assurance required by the financial function. As an arrangement approaches systemic importance, evidence should mature from functional demonstration to component performance, representative end-to-end workload testing, stress and degradation testing, production evidence and continuing reassessment.

Performance evidence should distinguish what was measured from what is inferred. As reliance grows, functional demonstrations should mature into end-to-end testing, production evidence, stressed evidence and continuing assurance.

⁹ Bank of England, Sterling-denominated systemic stablecoins: Policy Statement and consultation on draft Code of Practice (22 June 2026).

¹⁰ Bank of England, Sterling-denominated systemic stablecoins: Policy Statement and consultation on draft Code of Practice (22 June 2026), section 2.2.3.

IV. Observed Behaviour and Demonstrated Operational Evidence

Supervisory relevance. An accurate architectural description does not by itself establish how the complete arrangement behaves under the workload, dependencies and failure conditions relevant to the financial function. A systemic suitability judgement should therefore be supported by evidence of observed behaviour as well as intended design.

Relevant consultation questions: Q2, Q3 and Q4

Systemic financial infrastructure should ultimately be assessed according to how it actually behaves, not solely according to how its architecture is described.

Technical documentation, white papers, benchmark reports, governance frameworks and laboratory demonstrations can all provide useful information. They explain intended design and assumptions. For prudential purposes, intended design and observed behaviour should remain distinct.

For prudential purposes, the relevant objective is an evidence-based understanding of how the infrastructure performs the functions upon which the systemic financial outcome depends.

This response therefore distinguishes intended design, observed behaviour and the evidence available to support a supervisory conclusion.

IV-A. Declared Capability and Demonstrated Behaviour

- Fast confirmation does not establish what constitutes settlement finality.
- High throughput does not establish sustained successful finalised throughput under representative workload.
- Decentralisation does not identify the practical control surface.
- Immutability does not establish that recorded state is correct, complete or authoritative.
- Twenty-four-hour technical availability does not establish the ability to process stressed redemption demand continuously.

The distinction is analogous to conventional prudential practice. Supervisors do not infer operational resilience from an architecture diagram alone. They rely upon testing, operating records, incident experience, governance evidence, recovery exercises and observed performance.

IV-B. Supervisory Observability

Operational characteristics that cannot be reliably observed are difficult to supervise. This response uses 'supervisory observability' descriptively to mean the ability of supervisors and regulated entities to obtain reliable evidence concerning the operational facts upon which prudential conclusions depend; it is not proposed as a separate supervisory regime.

Supervisory observability does not imply continuous transmission of all operational data to the supervisor. It means that material prudential facts are defined, retained, reproducible and capable of being produced within a timeframe appropriate to the supervisory decision.

For a material operational assertion, observability should preserve not only the reported value but also the proposition being asserted, the producer or source of the underlying state, the observer or reporting party, the authority of that source for the proposition, the observation time, the evidence supporting it and any remaining dependency or condition that limits what can safely be inferred. A producer assertion, an observer-derived view, a cryptographic or reference artifact and a business recognition event can all be useful evidence without being equivalent propositions.

Observation also has a timing dimension. For a material proposition, the arrangement should distinguish event time, observation time, report time and reliance or action time. A representation can have been accurate when observed yet become stale, provisional or insufficient by the time an institution acts. Likewise, a query can improve visibility without itself establishing durability, finality, authority, business completion or that the returned state satisfies the assurance threshold required for the next financial action.

IV-B.1. From Observed Behaviour to Prudential Confidence

The concepts in this response form a deliberate progression. Observed operating behaviour records what the system does in practice. Supervisory observability concerns the ability of institutions and supervisors to establish the material facts when needed. Continuing operational evidence supports justified confidence that the relevant prudential outcomes remain satisfied as the system operates and changes.

The objective is not observability for its own sake. The objective is prudential confidence: evidence-based confidence in the properties upon which the regulated financial function and, at systemic scale, financial stability depend.

The relevant evidence can include settlement progression, governance decisions, dependencies, recovery, capacity limits and integrity controls that can be demonstrated rather than inferred.

- Which transactions are currently pending, failed, duplicated or retried?
- Which state transition is authoritative, and what is the relevant ordering?
- When did finality occur, and what can still alter the state?
- Which infrastructure dependency is degraded?
- Which privileged actions occurred, by whom and under what authority?
- What capacity remains available and what recovery process is active?
- Can the system's state at a specified historical time be reconstructed?

The absence of evidence can itself become a source of uncertainty. If institutions cannot determine quickly whether a transfer is pending, failed, finalised or subsequently altered, uncertainty about infrastructure can become uncertainty about ownership, liquidity and redemption.

IV-C. Point-in-Time State Reconstruction

A systemic arrangement should be capable of answering a deceptively simple question: what was true at a specified point in time?

That question becomes important after network disruption, emergency governance action, a sequencer restart, bridge failure, software upgrade, wallet or MPC incident, duplicate or retried submissions, freeze action or disagreement between issuer books and transaction infrastructure.

Point-in-time reconstruction is particularly important for regulator-facing assurance and should be considered alongside finality evidence, failure and retry records, operator maps and privileged-action logs.

The reconstruction test should focus on what could reasonably have been known at the decision point, not only on the eventual outcome. That distinction matters where a participant must decide whether to retry, replace, pause, release liquidity, recognise a redemption or resume service while sources disagree or the authoritative outcome is not yet established.

IV-D. Demonstrated Operating Behaviour Should Precede Critical Reliance

The purpose is not to eliminate technological uncertainty. It is to reduce avoidable uncertainty about the properties on which systemic reliance is placed. As stablecoins move further into payments, treasury, collateral and settlement workflows, material properties should be understood before they become critical dependencies, under representative and stressed conditions, and after material architectural change.

V. Capacity, Measurement Provenance and Production Fitness

Supervisory relevance. The relevant capacity measure is the workload that the complete arrangement can sustain at the required level of assurance, including under stress, rather than the number of technical actions that an individual component can process in isolation.

Relevant consultation questions: Q2, Q3 and Q4

Capacity is one of the most discussed characteristics of distributed infrastructure and one of the easiest to misinterpret. More than a decade of experimentation has produced substantial evidence of functional feasibility: issuance, transfers, atomic workflows, programmable business logic, tokenised assets and settlement models can work across many architectures.

As contemplated use moves from experimentation toward systemic financial activity, however, the evidentiary question changes.

Does the evidence establish only that the workflow can function, or does it establish that the complete arrangement can sustain the contemplated financial workload under representative production and stress conditions?

Total network capacity is not necessarily the capacity available to a systemic financial function. Where infrastructure is shared, multiple users and workloads draw on the same available processing capacity. During periods of stress, ordinary competing activity can coincide with protocol and system activity, failed or delayed attempts, retries and other infrastructure-generated workload, reducing the capacity available to the regulated financial function. Capacity evidence should therefore address the sustainable successful workload available to that function at the required assurance level under plausible concurrent demand, rather than relying only on peak network throughput or isolated benchmark results.

V-A. Functional Demonstration Is Not Production Validation

A controlled experiment can demonstrate that a token can be issued, several institutions can exchange it, a contract can execute, or a redemption workflow can complete. That proves something important: functional plausibility.

It does not necessarily establish sustained system capacity, concurrency behaviour, realistic geographic latency, failure rates, retry amplification, persistence, data-availability constraints, sequencing bottlenecks, stressed finality, recovery, operational support or redemption-effective capacity.

Functional feasibility, component performance, end-to-end performance and production readiness are related but distinct forms of evidence.

The policy context is also moving from experimentation toward persistent market-scale use. In their September 2026 wholesale-tokenisation Feedback Statement, the FCA and Bank report respondent pressure to move beyond pilots and sandboxes toward full production, scale and permanence, and state that they accept the need for scale and permanence. That development reinforces the need to distinguish functional demonstration from evidence of production fitness.¹¹

V-B. Measurement Provenance

Material performance claims used to support systemic suitability should preserve measurement provenance.

- what exact component or lifecycle was measured;
- the workload, transaction complexity, topology and geography;
- participant count, hardware and duration;
- which dependencies were included or excluded;
- the finality criterion applied;
- whether failed transactions were counted as throughput;
- whether retries were counted as new transactions;
- **behaviour at saturation and during recovery;**
- and the conclusion being inferred from the result.

A production-readiness assessment should identify what a stated result actually measured: for example, signature verification, message propagation, sequencer intake, local execution, proof generation, database writes or finalised state transfer. It should also identify whether wallet, custody, RPC, consensus, data availability, bridge and redemption workflows were included.

¹¹ Financial Conduct Authority and Bank of England, FS26/1: Tokenisation in Wholesale Financial Markets (14 September 2026), pp. 11-12.

The conclusion drawn from a performance test should not exceed what the test actually measured.

FROM COMPONENT MEASUREMENT TO FULL-LIFECYCLE EVIDENCE

The measured scope and completion criterion determine what a performance result can support.




Figure 4. Measurement Scope and Evidentiary Strength

V-C. Institutional Engineering Capability and Production Assurance

The distinction between functional demonstration and production validation raises a related institutional question: whether firms considering distributed infrastructure for material financial activity possess sufficient engineering capability to independently evaluate the evidence on which production-suitability conclusions depend.

Expertise in tokenization, application development, smart-contract implementation, compliance and business-process integration does not necessarily provide the systems-engineering expertise needed to evaluate the underlying transaction and state-transition infrastructure where that infrastructure is material. As reliance increases, institutions should be capable of distinguishing architectural claims and component-level measurements from evidence concerning the behaviour of the complete system under conditions representative of the contemplated financial workload.

Independent assurance requires the institution to understand what was actually measured, what assumptions were made, and whether the resulting evidence supports the production conclusion being drawn. Component-level evidence should not be used to support a system-level conclusion unless the relationship between the measurement and that conclusion is demonstrated.

The objective need not be a prescribed professional certification. A more immediate priority is a common engineering-assurance discipline: sufficiently consistent terminology, measurement methodology, institutional competency and production-readiness criteria to distinguish what a system is designed to do, whether a workflow can execute on it, and whether the complete system has been demonstrated to operate that workflow safely at the required production scale.

V-D. Gross Activity Is Not Necessarily Successful Economic Throughput

Under stress, unsuccessful submissions, timeouts, retries, duplicates, partial completion and delayed finality all consume capacity without necessarily representing unique economic activity successfully completed.

A decision-useful framework should distinguish gross attempts from successful execution, successful finalised economic transactions and complete regulated-lifecycle outcomes.

Retry behaviour matters because it can create nonlinear load. Degradation can produce failures; failures produce retries; retries increase offered load; increased load can create further degradation.

V-E. Sustained Capacity and Production Evidence

Systemic infrastructure should be assessed for the workload it may need to carry for the duration of a stress event, not only for a short peak. Relevant measures include sustained successful finalised throughput, latency distributions, completion, failures, retries, backlog, degradation thresholds, finality latency and recovery.

For a systemic arrangement, capacity is therefore not merely a performance characteristic. It is also a continuity property. The relevant question is whether the system remains within defined operating bounds as load increases, degrades in a known and observable manner at and beyond saturation, and can recover without creating ambiguity about transaction state or financial obligations.

A useful measurement ladder is therefore: theoretical or claimed capacity; observed gross activity; successful activity; activity reaching the required assurance or finality threshold; and complete business outcomes. Degraded-state capacity should be considered separately: the sustained successful or assured workload that remains possible while latency, failures, dependencies or recovery conditions deteriorate. Useful capacity is workload sustained at the assurance level the financial function actually requires.

As systemic significance increases, controlled testing should increasingly be supplemented by production evidence. Production history is valuable, but it should not be treated as dispositive where historical workloads differ materially from the financial workload being contemplated.

V-F. Empirical Network Evidence

Appendix A sets out a conservative evidentiary protocol and selected illustrative evidence. The purpose is not to rank networks or convert observed activity into a universal capacity claim. It is to demonstrate why advertised, benchmark, peak, observed and sustained production throughput are different measurements, and why a systemic suitability conclusion should not exceed the evidence available for the specific architecture, workload and regulated lifecycle being assessed.

VI. Settlement, Finality, Reversibility and Authoritative State

Supervisory relevance. Where a transfer is described as confirmed or final, the relevant questions are what proposition has been established, by which system or authority, and what conditions or dependencies remain before the financial obligation is complete.

Relevant consultation questions: Q1, Q2, Q3 and Q4

Finality is frequently used as though it describes one state. In practice, execution, confirmation, technical finality, issuer recognition, settlement finality and legal consequence can be distinct. Their ordering, coincidence and legal significance depend upon the architecture and applicable legal framework; none should be inferred solely from a technical label.

CPMI-IOSCO guidance places settlement finality at the centre of the PFMI analysis for systemically important stablecoin arrangements. The practical challenge for newer architectures is demonstrating exactly where and when that outcome occurs across the complete lifecycle.¹²

VI-A. Confirmation Is Not Necessarily Finality

A confirmation may show that a transaction has progressed. It does not necessarily establish that state can no longer change. Depending upon architecture, confirmed state may remain subject to reorganisation, challenge, proof completion, underlying-layer settlement, administrative intervention, bridge completion or another mechanism.

What exact state is being relied upon, and what, if anything, remains capable of changing that state?

VI-B. Finality and Reversibility Are Different Questions

A payment can be final and still be followed by a refund or compensating transaction. Conversely, a transaction represented as complete may remain technically capable of reversal or modification. Reversal of original state should therefore be distinguished from a subsequent refund, dispute payment, indemnification or compensating transaction.

VI-C. Layered Finality

Layered infrastructure can separate rapid execution from ultimate settlement. Execution layers, challenge periods, proof-generation stages, data-availability layers, bridges and underlying settlement layers may each introduce a different state or latency.

The presence of layers is not itself a deficiency. Supervisors should understand which layer produces the state used by participants, which layer determines ultimate finality, the latency between them, which mechanisms can alter earlier state and what economic exposure exists during the interval.

VI-D. Authoritative State

Potentially inconsistent records can exist across a transaction network, an L2 environment, bridge, wallet, custodian, exchange, issuer ledger, redemption service, indexer, RPC provider or bank settlement system.

Which record or determinable system state is authoritative for the relevant regulated obligation?

A component can be technically final while the regulated financial obligation remains incomplete. Source-layer finality, bridge or cross-domain evidence, issuer recognition, bank cash settlement and other business states should therefore be treated as separately established propositions where the arrangement requires them.

The answer may differ by stage. What matters is that the relationships are specified, observable and reconcilable.

VI-E. Timestamps and Ordering

During stress, timestamps and ordering can determine which redemption arrived first, whether a transfer preceded a freeze, whether an issuer had received a coin before an operational cutoff, or which institution carried exposure at the point of failure.

A systemic arrangement should therefore provide sufficient event evidence to reconstruct what happened, in what order, when, and according to which authoritative source.

VI-F. Finality Under Stress

Finality models should be evaluated under degraded conditions, including whether latency rises with load, whether the elapsed time to satisfy challenge, proof, publication or underlying-settlement conditions increases, whether congestion affects bridge completion and whether downstream systems can distinguish provisional from final state.

Median finality during ordinary operation is not sufficient evidence for a systemic stress case. Tail and stressed conditions matter.

¹² CPMI-IOSCO, Application of the Principles for Financial Market Infrastructures to stablecoin arrangements (13 July 2022).

VI-G. Technical State and Legal Consequence

A technical infrastructure does not determine legal finality merely by labeling a transaction final. Likewise, a legal rule cannot eliminate the technical dependencies necessary to establish the facts to which the rule applies.

Execution → Confirmation → Technical Finality → Issuer Recognition → Settlement Finality → Legal Consequence

These are distinct propositions rather than a universal sequence. Depending on the architecture and legal framework, some may coincide, occur in a different order, or be established by different systems or authorities.

The distinction is also visible in the broader UK tokenisation policy debate. In FS26/1, respondents argued that contractual finality alone was insufficient for certain blockchain-settled transactions because statutory settlement finality provides insolvency protection against third parties; the FCA and Bank recognised the importance of settlement and wider legal certainty. That is a legal-policy question distinct from technical finality, but it illustrates why the two propositions should not be collapsed.¹³

¹³ Financial Conduct Authority and Bank of England, FS26/1: Tokenisation in Wholesale Financial Markets (14 September 2026), p. 13.

VII. Governance, Control and Effective Decentralisation

Supervisory relevance. The control assessment should identify who or what can change the financial state, the rules governing that state, or participants' ability to act upon it.

Relevant consultation questions: Q2, Q3 and Q4

Governance becomes prudentially relevant because a systemic monetary claim can depend on actors or mechanisms with power to change outcomes even where day-to-day processing appears automated or distributed. The supervisory task is not to prefer decentralised or centralised governance. It is to identify the effective control surface and understand which powers can affect issuance, transfer, settlement, redemption, access or recovery.

VII-A. Nominal Decentralisation Is Not an Operational Control Map

This is not an argument that more decentralisation is always safer. Concentrated authority can make accountability and emergency action clearer; distributed authority can reduce some single points of control. The relevant issue is whether the actual powers, thresholds, dependencies and recovery mechanisms are known and compatible with the financial function.

Validator counts, stake concentration, node distribution, client diversity and Nakamoto-type coefficients can all contribute useful information. No single measure establishes the practical control surface of a systemic arrangement.

The key question is not merely how many participants exist. It is: who can change outcomes?

- validators or block producers;
- sequencers;
- protocol developers and foundations;
- smart-contract administrators;
- bridge and relayer operators;
- keyholders and MPC administrators;
- RPC and infrastructure providers;
- software-client maintainers;
- cloud or hosting providers;
- emergency or recovery committees.

Control mapping should include material powers to: pause, upgrade, censor, reorder, prioritize, replay, blacklist or recover, together with the entities and mechanisms through which those powers can be exercised.

Control surface	What it can affect	Supervisory question
Keys / administrators	Pause, freeze, mint/burn, permissions, upgrades	Who holds authority, under what threshold, with what evidence and recovery?
Ordering / sequencing	Priority, delay, exclusion, ordering	Can ordering alter a regulated outcome or access to scarce capacity?
Consensus / validation	Accepted state and conflict resolution	What establishes authoritative state and under what fault assumptions?
Software / upgrades	Rules by which valid state can change	Who can change the rules, how quickly, and what triggers reassessment?
Bridge / cross-domain	Movement and state recognition between domains	What is final on each side, and how is disagreement resolved?
Access / observation	Practical submission and view of state	For which bounded proposition, if any, is the source authoritative, and can the observation become stale, concentrated or unavailable?
Orchestration / recovery	Routing, timeout, retry, failover, emergency action	Can workflow logic create a second economic action before the first is resolved?

Figure 5. Effective control surface: the relevant question is who or what can materially change the financial outcome.

VII-B. Decisive Participation Sets

TauriPay proposes identifying the smallest practical combination of actors, authorities or dependencies capable of materially affecting a relevant outcome. Different decisive sets may exist for consensus, ordering, censorship, upgrades, emergency pauses, bridge operation, key recovery, minting or burning, and transaction access.

The objective is not a new universal decentralisation score. It is to identify practical concentration where it exists.

VII-C. Affiliation and Common Incentives

Raw operator counts can obscure common ownership, delegated control, commercial affiliation, governance participation or shared infrastructure. Apparent diversity may therefore overstate actual independence.

Affiliation can matter particularly where infrastructure participants also perform market-making, custody, exchange, distribution, liquidity or bridge functions, because correlated operational and economic incentives can arise.

VII-D. Ordering Authority Can Become Financially Consequential

During stress, ordering may determine which redemption is processed first, which transfer precedes a freeze, which participant obtains scarce capacity or whether fee or sequencing policy changes economic outcomes.

Different architectures can use arrival time, fee priority, proposer discretion, sequencer policy, deterministic rules, causal relationships or application logic. The mechanism need not be identical across systems, but it should be understood and evidenced where ordering affects a regulated outcome.

In some distributed architectures, transaction visibility and ordering authority can also create economically consequential behaviour that can take forms different from those typically encountered in conventional payment or settlement infrastructure. Examples can include selective inclusion or exclusion, fee- or auction-based priority, proposer or sequencer discretion, and forms of value extraction associated with pending-transaction visibility, often discussed under the broad heading of maximal or miner extractable value (MEV). The prudential issue is not the label itself; it is whether ordering and inclusion can alter the financial outcome, who can exercise that influence, and whether the effect is observable, bounded and appropriately controlled.

VII-E. Privileged Authority

Supervisory relevance. A financial instrument can operate on distributed infrastructure while particular actors retain powers capable of changing the conditions under which it transfers, settles or recovers. Those powers may be legitimate and may improve safety. Where an administrator, keyholder, committee or upgrade process can pause activity, freeze or blacklist assets, change software or permissions, mint or burn, replace contracts or coordinate recovery, that authority forms part of the effective control structure on which the financial function depends.

For systemic use, the relevant question is therefore not whether privileged authority exists, but whether its scope, governance and consequences are understood. An emergency key capable of protecting users during an exploit can be a resilience mechanism. The same authority can become a concentration, availability or state-integrity dependency if its exercise, compromise or unavailability can materially alter the financial obligation.

The control map should distinguish ordinary operating authority from exceptional or privileged authority and should identify the conditions under which each can be exercised.

For each material privileged control, supervisors should understand who possesses authority, how authorisation occurs, whether thresholds apply, what state the action can alter, what evidence is produced, how the control is tested and how misuse or compromise is detected and recovered from.

VII-F. The Same Questions Across Architectures

Permission changes the distribution of authority; it does not eliminate the need to understand authority. A tightly controlled private system may make accountability and intervention paths easier to identify while concentrating more power in fewer actors. A public system may distribute some powers more broadly while leaving the issuer dependent on governance, ordering, access or finality mechanisms it does not control. The prudential inquiry should follow the actual powers and dependencies rather than infer them from the architecture label.

These questions apply to private and conventional infrastructures as well as public networks. A private system can present concentrated operator, vendor, administrator and cloud risk. A public system can provide production diversity that some private systems lack.

Technology neutrality should produce consistency of prudential inquiry, not an assumption of operational equivalence.

VIII. Third-Party Dependencies, Operational Resilience, Recovery and Runbooks

Supervisory relevance. Legal accountability can remain with the issuer even when facts needed to perform the obligation are produced, observed or transmitted through systems outside the issuer's control. The assessment should therefore consider not only the resilience of critical service providers, but whether the complete dependency path preserves a coherent and actionable financial state.

Relevant consultation questions: Q1, Q2, Q3 and Q4

A systemic stablecoin issuer may remain legally accountable for an obligation while relying operationally upon systems it does not directly control.

VIII-A. Dependency Mapping

A useful dependency inventory should connect each material service to the regulated outcome it supports:

Dependency → Function → Regulated Outcome → Failure Effect → Failover → Recovery

This is more useful than a vendor list because it identifies why the dependency matters.

VIII-B. Bounded Authority and the Complete Dependency and Assurance Path

Third-party risk and dependency-path assurance are related but distinct. Third-party governance asks whether a provider is appropriately selected, controlled, resilient and substitutable. Dependency-path assurance asks what financial proposition is supplied through that provider, whether that proposition is authoritative for the next action, how it is observed, what evidence supports it and what remains unresolved.

For each material component or boundary, a useful supervisory sequence is:

**Proposition → Authority → Observation → Evidence → Durability / Indeterminacy → Remaining
Dependency → Safe Next Action**

This analysis is architecture-neutral. Relevant components can include custody and key management; access and RPC infrastructure; orchestration and routing; application logic; sequencers, validators and base ledgers; data-availability, proof, bridge and cross-domain mechanisms; cloud and software dependencies; and conventional payment or settlement components.

Each component can be authoritative for a bounded proposition without being authoritative for the complete business outcome. Custody can establish authorisation without destination execution; a gateway can establish receipt or routing without outcome; an observer can report state without being authoritative for it; and an orchestrator can determine workflow action while relying on other systems for execution or settlement facts. End-to-end resilience should therefore not be inferred solely from the resilience of individual components.

Where a gateway, orchestrator or other intermediary reports downstream status, the assessment should distinguish the proposition the intermediary establishes within its own authority from a proposition established by another domain that the intermediary observes, derives or represents. A representation of L1, L2, bridge or other destination state should therefore preserve the source and authority for the underlying proposition, together with the evidence and assurance conditions on which the representation relies. An intermediary may separately assume or stand behind a business obligation, but that does not by itself make it authoritative for the underlying destination state.

Operational resilience and structural/adversarial resilience are related but distinct. Operational resilience asks whether the arrangement can continue, degrade within understood bounds and recover through disruption. Structural/adversarial resilience asks whether the mechanisms on which the financial function relies continue to preserve the required state, authority, ordering and assurance properties when participants, messages, keys, operators or dependencies behave incorrectly or maliciously. The relevant evidence should follow the properties being relied upon rather than assume that availability alone establishes resilience.

VIII-C. State Coherence and Disconnection Risk

A composite system can remain technically available while losing alignment around the financial proposition on which the workflow depends. For purposes of this response, 'disconnection risk' describes this analytical condition: different components remain operational but their observations, timing, authority or assumptions about the relevant financial state diverge sufficiently to make the next action unsafe or the obligation difficult to reconstruct.

Illustrative forms include an execution domain reporting completion before a required assurance layer has reached the threshold relied upon by the financial workflow; an observer or indexer reporting a state that is stale or later changes; a bridge establishing source-side locking while destination-side completion remains unresolved; an issuer not yet recognizing a state reported by an external infrastructure provider; or a payment leg completing while the corresponding token-side obligation remains unreconciled.

The concern is not that every such divergence is a failure. Many are ordinary intermediate states. The prudential requirement is that the arrangement know which divergences are permissible, which source is authoritative for the proposition at issue, how long uncertainty can persist, what actions are prohibited while state is indeterminate, and how authoritative state is reconstructed when components disagree or become unavailable.

VIII-D. Orchestration Can Amplify Uncertainty

Orchestration deserves separate attention because it can determine economically consequential workflow actions without being authoritative for the underlying financial state. An orchestrator may determine routing, correlation, timeout handling, retry or replacement, failover and progression to the next workflow step while relying on another component for execution, finality, settlement or redemption facts.

That separation can become consequential under stress. If an observation times out while the underlying instruction remains unresolved, an automated retry or reroute can create a second economic action before the first action's fate is known. The resulting risk is not captured adequately by a simple availability metric. It depends on correlation between the original and subsequent attempts, authoritative-state determination, timeout semantics, retry policy and the ability to reconcile late completion.

TauriPay distinguishes known failure from indeterminate outcome because a controlled failure can enter a defined exception path, while an indeterminate outcome can make the next economic action itself a source of risk unless the workflow first establishes the relevant authoritative state.

VIII-E. Common-Mode Dependencies

Nominally independent providers may share a cloud provider, data centre, software client, bridge, RPC platform, key-management provider, underlying ledger or development team. Redundancy is not meaningful if the same underlying dependency can disable, compromise or materially distort both paths.

VIII-F. Runbooks Convert Design Into Operational Capability

Resilience is not demonstrated solely by architecture. Material scenarios should have tested runbooks that establish detection criteria, incident ownership, escalation, decision authority, communication, treatment of transaction state, reconciliation, recovery, reopening criteria and evidence preservation.

- capacity saturation and retry storms;
- failed submissions;
- MPC or key-management impairment;
- sequencer outage or validator disruption;
- bridge delay or failure;
- RPC/indexer divergence;
- selective access, censorship or conflicting cross-domain/observation evidence;
 - privileged governance intervention;
- compromise or misuse of privileged, administrative or key authority;
 - software-release failure;
 - mass redemption;
 - regulator evidence production.

These classes of disruption should be supported by tested runbooks appropriate to the architecture and financial function.

VIII-G. Release and Change Management

Protocol upgrades, client releases, smart-contract modifications, bridge changes, governance changes, security patches and configuration changes can alter properties upon which the original prudential conclusion depended.

Requirement → Validation → Deployment → Production Evidence → Material Change → Reassessment

Material changes can include changes to consensus or sequencing, bridge or interoperability architecture, governance or effective control, major software releases, emergency intervention mechanisms, material outsourcing or dependency arrangements, workload profile, or the finality model relied upon by the financial workflow. Where such a change could alter a relied-upon prudential property, the prior evidentiary conclusion should be reassessed rather than carried forward automatically.

Emergency changes deserve special attention because rapid intervention can alter ordering, validation, finality, account state, access or recovery assumptions.

VIII-H. Recovery Is More Than Restart

Technical restoration asks whether the service is running again. Financial-state recovery asks whether the arrangement can determine which transactions succeeded or failed, which were duplicated or retried, what state is authoritative, what obligations remain outstanding and whether institutional records reconcile.

For systemic money, recovery should include the ability to reconstruct and reconcile authoritative financial state.

A recovery time objective establishes when a function is expected to recover. It does not, by itself, establish the financial state to which the arrangement must recover. For a financial function dependent on distributed state, recovery assurance should therefore address both time and state: when the function must resume, and what authoritative financial state must be reconstructed or reconciled when it does.

Where snapshots, checkpoints or other recovery records are used, their existence alone does not establish recoverability of the financial state. The relevant question is whether state transitions occurring after the recorded recovery point but before disruption can be identified, replayed, rederived or otherwise reconstructed so that the restored arrangement can establish the state that had become authoritative for the relevant proposition before disruption. Any gap between the state on which participants or institutions were permitted to rely and the state that can be reconstructed after disruption should be identified, bounded and tested. For the supervisory assessment proposed here, recovery evidence should therefore demonstrate continuity of the financial state authoritative for the relevant proposition across the recovery boundary, rather than merely restart of the underlying infrastructure.

VIII-I. Resilience Evidence Should Include Exercises

The objective is not to prove that disruption is impossible. It is to demonstrate a repeatable operational loop:

Detect → Contain → Determine State → Recover → Reconcile → Evidence

VIII-J. Safety, Liveness and Deliberate Halt Behaviour

A distributed system can preserve safety by refusing to advance when it cannot establish a valid next state. In some circumstances, halting before committing an invalid or conflicting state is the correct engineering behaviour. For a regulated financial function, however, a safe halt remains an operational event: instructions can remain unresolved, queues can grow, reattempts can increase, dependent workflows can stop and liquidity or redemption obligations can continue while the infrastructure is unavailable.

Fit-for-purpose assessment should therefore distinguish safety from liveness and identify the conditions under which the system is expected to stop making progress. Relevant questions include what triggers a safety halt; whether the trigger is protocol-determined or operator-mediated; whether the system demonstrably halts before violating the relied-upon safety property; how pending and rejected instructions are treated; who can authorize or effect recovery; how authoritative state is established before restart; and what evidence demonstrates the behaviour under representative fault and, where material, adversarial conditions.

A claim that a system 'fails safely' should therefore be testable. The prudential question is not only whether incorrect state is prevented, but whether the financial consequences of the resulting loss of liveness are understood, observable, bounded and recoverable.

VIII-K. Substitutability, Migration and Re-Instantiation

Some material infrastructure dependencies can be substituted through ordinary failover. Others are embedded in the mechanism by which the instrument's operative state is established or changed. Replacing such a dependency may require migration of authoritative records, coordinated state transition, or re-instantiation of the instrument on a different environment.

For material dependencies, exit planning should therefore identify the conditions under which migration would be required; the portability and verifiability of authoritative records; the authority to initiate and complete migration; treatment of pending or disputed transactions; prevention of duplicate or conflicting state; and the criteria for resuming normal operation.

A synchronized copy of historical state can assist recovery, but it should not be assumed to provide an independent mechanism for producing valid future state transitions. Where substitution is operationally or legally constrained, that limitation is itself relevant to the dependency assessment.

IX. The Systemic Stress Pathway: When Infrastructure Stress Becomes Confidence Stress

Supervisory relevance. Infrastructure stress can become a financial-stability issue where technical delay, uncertainty or disconnection changes what holders and intermediaries can safely do with the monetary claim. The transmission mechanism is not limited to outage: uncertainty about state can itself contribute to liquidity demand, retries, precautionary behaviour and confidence effects.

Relevant consultation questions: Q2, Q3 and Q4

Many stablecoin risk analyses appropriately begin with a loss of confidence in the instrument and then considers redemptions, reserve liquidation, market effects and contagion. Systemic analysis should also consider the reverse pathway: can infrastructure stress itself become a catalyst or amplifier of confidence stress?

IX-A. The Two-Layer Stablecoin Crisis


TauriPay proposes analyzing systemic stress across two interacting layers.

The financial layer includes redemption demand, reserve liquidity, asset liquidation, peg pressure, market-maker withdrawal, exchange contagion, banking interaction and official-sector response.

The infrastructure layer includes capacity, latency, failures, retries, transaction ordering, finality, sequencer or validator health, data-availability and proof dependencies, bridges and cross-domain state, access and observation infrastructure, orchestration and routing, custody/key controls, issuer-recognition systems, payment and settlement dependencies, privileged controls, software changes, recovery and authoritative-state determination.

TWO-LAYER SYSTEMIC STABLECOIN STRESS MODEL

Infrastructure stress can amplify financial stress – and financial stress can increase infrastructure load



Stress-test question: can the arrangement establish what happened, preserve authoritative state and continue the regulated function?

Figure 6. Two-Layer Systemic Stablecoin Stress Model

A relevant stress scenario is one in which market confidence deteriorates because the ecosystem cannot establish ordering, finality, capacity, control authority, intermediary exposure, recovery or authoritative state when those facts become most important.

IX-B. A Non-Hostile Stress Scenario

1. Adoption raises ordinary transaction demand through payments, treasury movement, trading and redemptions.
2. Demand approaches the sustainable operating envelope and latency rises.
3. Failures and timeouts increase; some transactions remain unresolved.
4. Wallets and applications retry delayed or unsuccessful instructions, amplifying offered load.
5. Participants become uncertain about queue position, finality and actual balances.
6. Market makers and intermediaries become defensive and reduce provisional exposure.
7. Redemption demand rises as some holders move back into bank money.
8. Redemption traffic itself adds load to infrastructure already under stress.
9. Infrastructure uncertainty becomes confidence uncertainty: the question shifts from whether reserves exist to whether the instrument can actually be moved and redeemed when needed.

This is a resilience-testing scenario, not a prediction about any particular architecture or market outcome. It does not assume collapse. It identifies a pathway by which operational degradation could become financially consequential even without an initial reserve impairment or hostile attack, and asks whether the arrangement can absorb rather than amplify the resulting stress.

The same transmission mechanisms can also be triggered by hostile or manipulative events. Sections VII and VIII address the control, authority and structural-resilience questions relevant to those cases.

IX-C. Composite Stress Without a Single Catastrophic Failure

The Bank's operational-resilience framework already requires extreme-but-plausible scenario testing. For composite stablecoin infrastructure, those scenarios should include cases in which no single component has suffered a conventional outage but the systems carrying the obligation no longer agree on, or cannot establish with sufficient assurance, the state on which the next financial action depends. The state, authority and assurance analysis in this response provides a way to test whether the complete arrangement preserves a coherent and actionable financial state through degradation, uncertainty and recovery. TauriPay suggests the following illustrative composite stress scenarios for that purpose:¹⁴

Scenario 1 - Indeterminate execution and automated retry. An orchestrator or wallet loses timely observation of an instruction, treats the timeout as unsuccessful and retries or reroutes it. The original instruction later reaches the relied-upon state. Unless attempt identity and authoritative-state reconciliation are preserved, uncertainty can become duplicate or conflicting economic action.

Scenario 2 - Cross-layer assurance mismatch. An execution layer reports completion, but the financial workflow relies on an additional assurance condition at another layer or domain. A downstream redemption, transfer or accounting action advances before that condition is satisfied. No component need be malfunctioning; the failure is an incorrect inference across an assurance boundary.

Scenario 3 - Observation and business-state divergence. An indexer, RPC provider, bridge service or other observer reports a state that the issuer, custodian or settlement component has not recognised, or the reported state later changes. Participants can then hold different but individually plausible views of whether the financial obligation has completed.

Scenario 4 - Financial-leg disconnection. A bank or payment-system leg completes while the corresponding token-side or issuer-recognition state remains unresolved, or vice versa. The relevant risk is not merely service availability but whether the arrangement can prevent, identify and reconcile an economically inconsistent composite state.

IX-D. Retry Amplification

Retry behaviour creates a nonlinear relationship between economic demand and infrastructure workload. Degradation can produce failures; failures produce retries; retries increase load; increased load can produce further degradation.

The appropriate metric is therefore not merely original demand but retry-adjusted workload.

Transaction demand should also not be assumed to map one-for-one to infrastructure workload. One business obligation can generate multiple validation, ordering, execution, publication, proof, observation, relay, reconciliation or recovery actions. Under stress, the relevant operating load can therefore reflect both workflow expansion and retry amplification, even when underlying economic demand has not increased proportionately.

Observation demand can amplify in parallel. As uncertainty increases, participants may issue more status queries, switch observation providers and initiate reconciliation or failover precisely when RPC, indexer or other observation infrastructure is degrading. Delayed or stale observations can then increase uncertainty further. Stress testing should therefore consider query- and reconciliation-adjusted observation load as well as transaction retry load.

IX-E. Interim Exposure: Who Carries Risk Between States?

When technical completion, finality, issuer recognition and cash settlement occur at different times, someone may carry the resulting exposure. The prudential analysis should identify who that party is, whether the exposure is funded or collateralized, whether it is visible, and whether multiple intermediaries can accumulate similar exposures simultaneously.

Potential carriers include exchanges, custodians, payment processors, market makers, affiliates, banks, bridges and liquidity providers. The assessment should therefore consider provisional credit, internal netting and intermediary balance-sheet exposure where these can transmit or amplify infrastructure stress.

IX-F. Ordering Under Scarcity

Where only part of incoming workload can be processed promptly, ordering becomes economically consequential. Queue position, fee priority, sequencer policy and preferred access can affect who obtains liquidity first.

TauriPay's prudential point is that a redemption queue is not merely a technical queue when its order determines who exits first.

This is consistent with the Bank's proposed requirement that full redemption requests be processed using fair and objective criteria and, under the June 2026 proposal, on a rolling first-come, first-served basis. The additional infrastructure question is whether upstream ordering, congestion or access conditions can materially affect which holders are able to reach the point at which that protected queue begins.¹⁵

¹⁴ Bank of England, Operational resilience: recognised payment system operators and specified service providers, Supervisory Statement (12 June 2026), paras 3.19-3.21.

¹⁵ Bank of England, Sterling-denominated systemic stablecoins: Policy Statement and consultation on draft Code of Practice (22 June 2026), section 2.2.3.

IX-G. Financial Assistance and Infrastructure Assistance Solve Different Problems

Liquidity support can address a financially sound issuer's inability to monetize reserves quickly enough. It does not independently resolve an unavailable sequencer, congested ledger, bridge failure, uncertain finality, MPC outage, corrupted release, inaccessible transaction path or inconsistent authoritative state.

Financial resilience and infrastructure resilience are complementary. They are not substitutes.

IX-H. The Supervisory Objective Is Evidence During the Event

A stress regime should not only determine whether the system survives. It should determine whether institutions and supervisors can establish what is happening accurately enough and quickly enough to make sound decisions while it is happening.

A regulator-facing evidence packet could include outstanding supply, pending/completed/failed/retried redemptions, finalised throughput, retry-adjusted workload, finality evidence, operator maps, privileged actions, intermediary exposures and activated runbooks.

X. Systemic Recognition Should Expand the Assurance Perimeter

Supervisory relevance. The consequences of an infrastructure weakness increase as an arrangement becomes more widely used, interconnected or difficult to substitute. Systemic recognition should therefore increase the amount and quality of evidence supporting reliance, even where the underlying technology has not changed.

Relevant consultation questions: Q1, Q2, Q3 and Q4

Systemic recognition changes the significance of operational dependency. As stablecoin use expands into payments, treasury, settlement, collateral and other important functions, impairment of the underlying infrastructure can affect a broader set of institutions and markets.

As systemic importance increases, the evidentiary perimeter should expand to include material infrastructure dependencies required for the issuer to discharge its regulated obligations.

X-A. Regulated Accountability and Operational Dependency Should Not Diverge

This does not imply that an issuer must control every external component or that the Bank must supervise every technical provider directly. It means that where the issuer's regulated performance materially depends on a component it does not control, the dependency, evidence limits, contingency arrangements and residual risk should remain visible in the suitability assessment.

A systemic issuer may remain legally accountable while successful performance depends upon a ledger, sequencer, RPC access layer, custody platform, MPC process, bridge, smart contract or other service. Where the dependency is material, assurance should not stop at the corporate boundary.

X-B. Materiality Determines Depth of Assurance

The depth of evidence should increase with scale, systemic importance, interconnectedness, complexity, substitutability, dependency concentration and severity of outcome if the dependency fails.

X-C. Infrastructure Dependency Assessment

As an arrangement approaches systemic recognition, the issuer should be able to produce an Infrastructure Dependency Assessment identifying the complete lifecycle, material dependencies, operators/control structure, capacity evidence, finality semantics, failure/recovery mechanisms, common-mode dependencies, privileged controls, testing/production evidence and conditions requiring reassessment.

X-D. The Glidepath Into Systemic Importance

The evidence process should begin before formal systemic recognition if material reliance is already developing. Waiting until the system has become critical to discover whether relevant evidence exists defeats the purpose of forward-looking prudential supervision.

X-E. Fitness for Purpose Must Permit a Negative Conclusion

A credible fit-for-purpose assessment should be capable of concluding either that an arrangement is not presently fit for the contemplated function or scale, or that the available evidence is insufficient to establish such fitness.

Technology neutrality does not mean that every technology must ultimately be approved. It means that any conclusion should concern the evidenced suitability of the particular arrangement for the contemplated reliance and should arise from evidence concerning required outcomes rather than architectural preference.

XI. Supervisory Observability and Continuing Operational Evidence

Supervisory relevance. During stress, material facts should be capable of being established within the timeframe in which supervisory and operational decisions must be made, rather than only reconstructed after the event. The evidentiary objective is therefore to make the relevant state, authority, dependencies and recovery position decision-useful when they matter.

Relevant consultation questions: Q3 and Q4

Periodic reporting, examinations, audits, governance reviews, stress testing and direct supervisory engagement remain essential. Increasingly software-defined infrastructure also creates the possibility of more timely and granular operational evidence.

The relevant objective is continuing operational evidence: the ability to support supervisory confidence over time with proportionate, independently verifiable evidence as the arrangement operates and changes.

Material operational properties demonstrated before deployment should remain capable of being evidenced after deployment and as the infrastructure changes.

XI-A. An Evidence Ladder

Declarative evidence. Policies, specifications, architecture documents, governance statements and provider representations establish intended design.

Demonstrated evidence. Tests, pilots, benchmarks, failover exercises and controlled demonstrations establish behaviour under specified conditions.

Operational evidence. Production completion, failures, retries, latency, congestion, outages, recovery and governance actions establish how the system has actually behaved.

Continuing evidence. Production metrics, recovery exercises, material-change testing, dependency monitoring, governance-action records and incident evidence establish whether important properties remain true.

Declared → Demonstrated → Operational → Continuing

XI-B. Continuing Evidence Does Not Mean Continuous Approval

Supervisors need not approve every software release, configuration change, governance action or transaction. Continuing evidence means that when a material prudential question arises, reliable evidence exists to answer it.

XI-C. Systemic Stablecoin Infrastructure Evidence Package

1. Lifecycle and material dependency map.
2. Settlement, certainty-of-outcome and finality definitions for each material layer.
3. Capacity evidence, including successful finalised throughput, failures, retries, retry-adjusted workload, latency and redemption-effective capacity.
4. Control and affiliation map covering material operators, keyholders, governance actors, providers and software dependencies.
5. Privileged-control records covering pauses, upgrades, blacklists, mint/burn actions and emergency interventions.
6. Dependency and concentration evidence, including common-mode and failover analysis.
7. Recovery evidence, including state reconstruction, reconciliation, pending/duplicate/retry treatment and reopening criteria.
8. Material-change evidence showing changes capable of affecting properties on which systemic suitability depends.
9. Stress-test evidence, including capacity stress, mass redemption, dependency outage, degraded finality and, where material, adversarial-control scenarios.
10. Measurement and observation methodology and provenance sufficient to determine what was measured or observed, when, from which source and what conclusion is being inferred.

XI-D. Evidence Should Be Available During Stress

Timing matters. An after-action report may help prevent recurrence but may not help the supervisor decide what to do while stress is unfolding. The regime should therefore support production of a defined evidence packet within a decision-useful period during material stress.

XI-E. Independent Verifiability

Important claims should, where practical, be supported by independently reproducible measurements, corroborating data sources, protected audit records, reconciled institutional records or other evidence that does not rely exclusively on a single narrative representation.

No single assurance method should become a proxy for production suitability. Source-code availability can strengthen inspectability and independent review, but fully open source code is not the only way to establish a reasonable evidentiary basis. Depending upon the function and risk, assurance may also be supported through independent technical review, controlled testing, formal analysis, protected audit access, execution and operational records, regulator access, reproducible measurements, attestations or other reliable methods. The relevant question is whether the evidence is sufficient for the proposition and level of reliance at issue.

Labels such as 'widely used,' 'battle-tested,' or 'institutionally supported' can provide context, but they are not substitutes for evidence of the specific properties on which systemic suitability depends. Where the property is material, the supporting methodology, assumptions, test conditions and results should be sufficiently documented to permit independent review, corroboration or re-performance proportionate to the risk.

XI-F. Material Change Should Trigger Reassessment

A systemic suitability conclusion should remain subject to reassessment where material change affects a property on which that conclusion depends. Systemic assurance should not remain permanent if consensus, sequencing, finality semantics, bridges, software, governance, keys, dependencies, capacity or recovery mechanisms change materially.

XII. Proportionality, Innovation and Technology Neutrality

Supervisory relevance. Architecture-aware supervision need not favour a particular technology. Different designs may satisfy the same prudential outcome, but the evidence supporting reliance should reflect the control, finality, capacity, recovery and dependency characteristics of the arrangement actually used.

Relevant consultation questions: Q2, Q3 and Q4

An infrastructure-evidence framework should not become a barrier to experimentation. The burden should scale with criticality.

XII-A. Requirements Should Scale With Reliance

A pilot should not be treated like systemic infrastructure. The evidentiary burden should increase with scale, criticality, interconnectedness, substitutability, complexity and potential impact of failure.

XII-B. Technology Neutrality Applies to Outcomes and Evidence

A technology-neutral framework should ask the same core questions of centralised systems, public DLT, private DLT, L1, L2, hybrid systems and future architectures: what state is authoritative, when is settlement final, what is sustainable capacity, what happens under stress, who can change outcomes, what dependencies exist, how does recovery work, and can the evidence be corroborated?

This suggests a fit-for-purpose discipline rather than an architecture ranking. The relevant conclusion is not whether a ledger, Layer 2, bridge, permissioned network or other component is good or bad in the abstract. It is whether the complete arrangement has been demonstrated, with evidence proportionate to the contemplated reliance, to satisfy the capacity, state, finality, ordering, resilience, recovery, governance and control requirements of the particular regulated financial function.

A technology-neutral framework can therefore be architecture-aware. Different technical arrangements create different commit paths, control surfaces, dependencies and failure modes; the evidentiary questions should follow those differences without prescribing a preferred architecture.

The FCA and Bank's September 2026 tokenisation Feedback Statement reinforces the same distinction: the authorities remain technology agnostic and apply the principle of 'same risk, same regulatory outcome,' while stating that equivalent treatment of tokenised and non-tokenised assets depends on the risks being comparable and sufficiently mitigated. Technology neutrality therefore does not require an assumption that different architectures have equivalent operating characteristics; it requires consistent prudential outcomes supported by evidence appropriate to the actual risk.¹⁶

XII-C. Public and Private Infrastructure Face the Same Evidentiary Logic

A public network should receive credit for production history, availability, broad participation and other demonstrated strengths where they exist. It should also be assessed for the specific finality, congestion, effective-control, governance, dependency and workload characteristics relevant to the contemplated function.

A private system should receive credit for explicit governance and operational control where demonstrated, but should not receive an automatic presumption of production fitness merely because regulated institutions participated in a pilot.

Conversely, a public network with substantial production history may provide stronger evidence concerning real congestion, failure, recovery and adversarial operating conditions than a newly deployed permissioned system supported principally by pilots or laboratory tests. Production evidence should receive evidentiary weight regardless of architecture; institutional sponsorship or permissioned governance should not substitute for it.

A systemic suitability conclusion should not exceed what the available evidence demonstrates.

XII-D. Evidence Can Accelerate Innovation

Clear evidentiary expectations can reduce uncertainty for developers and institutions. If the market knows that systemic suitability will require evidence concerning finality, capacity, failures, retries, control, recovery and dependencies, those properties can be engineered and instrumented earlier.

Design for observability before supervision requires it.

¹⁶ Financial Conduct Authority and Bank of England, FS26/1: Tokenisation in Wholesale Financial Markets (14 September 2026), pp. 8, 13.

XII-E. Preserve Outcomes as Technology Evolves

The durable part of the framework should be the required outcome rather than today's implementation mechanism. Require reliable settlement finality rather than a specific consensus method; sustainable capacity rather than a specific scaling architecture; accountable control rather than a particular governance form; and demonstrated recovery rather than a particular recovery design.

XII-F. The Proper Standard Is Not Perfection

No financial infrastructure is failure-free. The relevant question is whether material failure modes are understood, proportionately controlled, observable, recoverable and compatible with the systemic function the infrastructure is expected to perform.

XIII. Recommendations for the Current Code and Developing Supervisory Framework

Relevant consultation questions: Q1, Q2, Q3 and Q4

The recommendations are divided deliberately by regulatory home. Items 1-5 are limited clarifications relevant to the current systemic stablecoin Code and consultation. Items 6-19 are principally candidates for the supporting supervisory framework the Bank has indicated it will continue to develop, including operational resilience, third-party risk, reporting, record-keeping and application of the PFMI. Items 20-21 are broader supervisory and industry-assurance propositions that need not be embedded as prescriptive Code requirements. All should apply only where material to the regulated function and proportionate to systemic reliance.

Read as a supervisory programme rather than as 21 isolated controls, the recommendations address five linked capabilities: mapping the complete dependency path; defining state, authority and completion; establishing reliable measurement and production evidence; testing degradation, recovery and state reconstruction; and maintaining a fit-for-purpose conclusion as reliance and architecture change. The individual recommendations below identify where those capabilities can be embedded in the current Code, the developing supervisory framework or longer-term industry assurance work.

XIII-A. Current Code and Consultation

Regulatory home: current Code and June 2026 consultation. These recommendations clarify how existing outcomes such as redeemability, operational resilience, finality and issuer accountability should follow material infrastructure dependencies without prescribing architecture.

1. Clarify that continuous redeemability should be assessed across the complete redemption path where external infrastructure is material to holder delivery and issuer receipt.
2. Preserve an explicit issuer obligation to identify material infrastructure dependencies by function and consequence rather than by technology label or corporate boundary.
3. Ensure that operational resilience includes recovery of authoritative financial state, not only technical restoration of service.
4. Clarify that where regulated obligations depend upon technical finality, issuer recognition or settlement finality, the relevant state and its relationship to the applicable legal consequence should be defined and evidenced without assuming that technical and legal finality coincide.
5. Preserve the ability to require reassessment where material architectural or operational change alters a property on which systemic suitability depends.

XIII-B. Developing Supervisory Framework

Regulatory home: principally the Bank's developing 2027 supporting materials and supervisory practice, including operational resilience, third-party risk management, reporting and record-keeping, recovery planning and application of the PFMI. These recommendations are intentionally framed as matters the Bank could consider rather than immediate prescriptive Code rules.

6. Consider requiring an Infrastructure Dependency Assessment as systemic importance develops.

Measurement, evidence and production fitness

7. Promote common operational definitions for transaction attempt, execution, completion, finality, failure, indeterminate outcome, retry or replacement, throughput, sustainable capacity, assurance status or threshold and redemption completion.
8. Consider requiring measurement provenance for material capacity, latency and resilience claims.
9. Consider requiring representative end-to-end capacity evidence that distinguishes claimed or benchmark capacity, observed gross activity, successful activity, activity reaching the required assurance or finality threshold, degraded-state capacity and redemption-effective capacity where relevant.

State, authority and composite-path assurance

10. Consider requiring material dependency-path mapping that identifies, for each component or boundary relied upon by the systemic function, the proposition established, the relevant authority or source, the observation path, supporting evidence, remaining dependency and the action that may safely follow.
11. Consider requiring arrangements to define how authoritative financial state is determined and reconstructed when material components disagree, observations become stale, cross-domain completion is partial, or automated orchestration encounters an indeterminate outcome.

12. Consider incorporating infrastructure-generated amplification and material fault or adversarial-control conditions into existing extreme-but-plausible scenario testing where relevant, including indeterminate outcomes followed by retries, dependency degradation, cross-layer assurance mismatch and simultaneous redemption pressure.

Control, recovery and continuing evidence

13. Consider developing expectations for effective-control and affiliation mapping rather than relying on a single decentralisation metric.
14. Consider using existing operational-resilience testing and recovery expectations to require tested runbooks, where material, for failure classes involving authoritative-state reconstruction, reconciliation of unresolved or duplicate actions and safe reopening criteria.
15. Consider developing expectations for continuing production evidence and reassessment after material change.
16. Consider requiring a decision-useful evidence package during material stress, sufficient to establish the relevant state, who or what produced or asserted it, the supporting provenance, and what could reasonably have been known when material operational decisions were made.
17. Apply any such expectations proportionately and technology-neutrally.
18. Consider requiring proportionate exit and migration planning where a material infrastructure dependency cannot be substituted without movement or re-instantiation of authoritative state.
19. Consider developing proportionate assurance expectations for material technical claims, with the depth of specification, testing of relied-upon fault and adversarial assumptions, and independent review scaled to systemic importance and consequence of failure, without treating any single assurance mechanism, including source-code availability, as a substitute for fit-for-purpose evidence.

Longer-term assurance and explicit suitability conclusions

20. Encourage or recognise industry-led, architecture-neutral engineering-assurance methodologies that help institutions distinguish design claims, functional demonstrations and complete-system production evidence.
21. Consider expecting an explicit fit-for-purpose conclusion where materially relied-upon infrastructure supports a systemic function, with unresolved evidentiary gaps, conditions and required mitigants recorded before reliance increases.

XIII-C. Proposed Supervisory Principle

Before infrastructure becomes a material dependency of systemic money, the evidence supporting its fitness for purpose should be commensurate with the financial function, scale and stress conditions it may be required to support. That conclusion should follow the complete dependency and assurance path without requiring the Bank to certify a particular protocol or architecture. Where material change affects a property on which suitability depends, the evidentiary conclusion should be reassessed.

XIV. Conclusion

For systemic reliance, the Bank should be able to determine what state the financial outcome depends upon, who or what establishes or can change it, what evidence supports it, and whether the next action remains safe during stress and recovery.

Relevant consultation questions: Q1-Q4

A systemic suitability conclusion should not exceed what the available evidence demonstrates.

Stablecoins can become part of systemic money. If they do, the prudential question cannot stop at backing and legal redeemability. The infrastructure through which the instrument is held, transferred, settled and redeemed can determine whether those protections are practically available when needed.

Backing is necessary. Redeemability is necessary. Systemic reliance also requires evidence that the complete arrangement can sustain the required workload, establish the state and finality on which institutions rely, distinguish known failure from indeterminate outcome, and recover and reconcile under stress.

Confidence should therefore be an output of assurance, not an input to it. Adoption, institutional sponsorship, transaction volume or years of operation may provide useful context, but they should not substitute for evidence that the properties required of a particular systemic financial function have been demonstrated. As reliance increases, that evidence should mature from design claims and functional demonstrations to measured production behaviour, stressed performance, demonstrated recovery and continuing reassessment after material change.

This is not an argument for or against distributed infrastructure. Public, private, permissioned, layered and conventional arrangements should be assessed by what they can demonstrate. The Bank need not certify protocols or eliminate technical risk. It should, however, be able to identify which infrastructure properties and dependencies are material to the regulated outcome, what evidence supports reliance upon them, what conditions or gaps remain, and what changes would require reassessment.

Those questions should be answered before the infrastructure becomes critical, not discovered only after the financial system already depends upon it. The developing framework should therefore build the supervisory capability to understand material authority boundaries, capacity, finality, failure behaviour, recovery and changing dependencies as systemic reliance grows. That is a continuing assurance task, not a one-time certification exercise.

The lasting supervisory question is not simply whether each component is functioning. It is whether the complete arrangement can establish, evidence and, when necessary, reconstruct the financial state on which the next safe action depends.

Respectfully submitted,

Mark Sulavka
TauriPay, a Public Benefit Corporation

Appendix A. Infrastructure Performance Evidence and Illustrative Measurements

Purpose. This appendix sets out the evidentiary protocol for quantitative and architecture-specific examples used in this response. It is intentionally conservative: a network-specific figure is included only where the source, measurement scope and filing-date relevance can be stated with sufficient precision for the narrow prudential inference drawn from it.

The Bank of England's DLT Innovation Challenge provides a useful illustration of the distinction. The final report examines settlement finality and security, scalability, network and asset control, and interoperability as interdependent design questions. It records hypothetical performance figures for certain combined approaches while stating that those outcomes involved trade-offs between speed, network distribution and fault tolerance and were not independently validated. The prudential point is narrow: a demonstrated design approach or hypothetical performance result is not, without more, a production-fitness conclusion. Source: Bank of England and BIS Innovation Hub London Centre, DLT Innovation Challenge 2025: Final Report, 12 May 2026, <https://www.bankofengland.co.uk/research/fintech/dlt-innovation-challenge-2025>.

A-1. Evidence classes

For capacity and activity claims, this response distinguishes five propositions: theoretical or claimed capacity; observed gross activity; successful activity; assured or finalised activity; and complete regulated or business outcomes. These should not be collapsed into one throughput number. Degraded-state capacity should be reported separately where material.

- Production activity: sustained successful economic throughput, workload composition and relevant production history.
- Completion quality: successful execution/finalisation, failures, unresolved transactions, retries and retry-adjusted workload.
- Finality: ordinary and stressed latency, relevant layer dependencies, reversibility/challenge conditions and issuer-recognition timing.
- Capacity: sustained rather than burst performance, saturation behaviour, backlog, degradation thresholds and recovery.
- Control and dependency: sequencers, validators, bridges, key/control functions, software/client concentration and common-mode dependencies.
- Measurement provenance: exact test scope, topology, geography, hardware, transaction complexity, duration, exclusions and finality criterion.

A-2. Inclusion test for an illustrative case

- The example must illuminate a proposition in the body rather than merely show that a named system experienced difficulty.
- Primary or technically authoritative sources should be preferred; secondary commentary should not carry a material factual claim by itself.
- The example must distinguish measured facts from inference and identify whether it concerns a component, controlled environment, complete lifecycle or production workload.
- The filing should avoid ranking networks or extrapolating from a single incident to an architecture class.
- Where evidence is ambiguous, stale or methodologically incomparable, the example should be omitted rather than forced into the paper.

A-2.1. Illustrative production incident: safety, liveness and recovery

Base reported two mainnet block-production outages on 25 and 26 June 2026, lasting approximately 116 minutes and 20 minutes. According to the network's published postmortem, a sequencer block-building bug allowed stale journal state to persist after a failed transaction; a subsequent valid transaction was then executed against that inaccurate state, producing a block with an invalid state transition that other nodes would not accept. The chain halted. Base reported that chain integrity was not compromised and funds remained safe, while new transactions were not included onchain and newly submitted transactions could be rejected as the transaction pool filled.

The example is relevant not because an outage establishes that the architecture is unsuitable, but because it illustrates the distinction between safety and liveness. Refusal to advance an invalid state can preserve an important safety property while simultaneously interrupting the regulated financial function. A fit-for-purpose assessment should therefore test both whether the system stops when it should and whether the financial consequences, recovery path and authoritative state are understood when it does.

Source: Base Engineering Team, 'Postmortem: June 25th Block Production Outage,' 27 June 2026, <https://blog.base.dev/postmortem-june-25th-block-production-outage>. The source reports two mainnet block-production outages on 25 and 26 June, lasting approximately 116 minutes and 20 minutes, and states that chain integrity was not compromised while transaction processing halted.

A-3. Measurement-scope taxonomy

- Communication or ordering benchmark: measures propagation, dissemination or ordering without necessarily including full admission, execution, state transition and finality.
- Partition or shard aggregation: derives aggregate capacity from multiple lanes or partitions; cross-partition coordination and shared-state effects must be separately demonstrated where relevant.
- Resource-budget arithmetic: converts gas, compute or other resource budgets into an estimated TPS using assumed transaction costs; this is a capacity estimate, not observed end-to-end throughput.
- Component execution benchmark: measures a client, database, proof system or execution path in isolation; component rate should not be presented as complete-system settlement capacity.
- Controlled end-to-end test: exercises the full transaction path in a defined environment; topology, geography, duration, state size, contention and fault conditions remain material to interpretation.
- Production observation: records actual live activity; observed demand is evidence of usage and behaviour, but is not by itself a measurement of maximum or stressed capacity.

A-4. Production outcome evidence

The production exhibit below uses reproducible public queries to illustrate why attempted activity, successful outcomes and application-level reverts should not be collapsed into a single throughput or failure measure. Its purpose is not to rank networks or imply direct comparability among heterogeneous outcome semantics.

For the exhibit, the relevant evidentiary record should preserve the observation period, transaction population, definition of successful and unsuccessful outcomes, whether reattempts can be identified, material exclusions, query methodology, extraction date and a public query or dashboard reference where available. Because public dashboards can update after filing preparation, any quantitative figure relied upon in the final submission should be retained as a dated snapshot or otherwise reproducible filing record.

Outcome semantics differ materially across architectures. A reverted application transaction, rejected submission, dropped transaction, timeout, consensus-layer failure and user reattempt are not interchangeable events. Any cross-network presentation should therefore preserve the native meaning of each dataset and limit the inference accordingly.

A-4.1. Cross-chain revert observations: heterogeneous outcome semantics

The cross-chain Dune series records transaction revert rates across multiple chains over the six-month period ending 16 September 2026. It should not be interpreted as a directly comparable infrastructure-failure ranking. A reverted application transaction can reflect contract or application logic rather than consensus or infrastructure failure, and the underlying transaction populations and outcome semantics differ across networks. Its value for this response is to illustrate why production activity should be interpreted using native outcome definitions, explicit numerator and denominator, a defined observation window and appropriate limits on cross-network inference.




Figure A-1. Transaction revert-rate observations across chains, six-month period ending 16 September 2026. Source: Dune query 8750094, visualisation 12777052, snapshot run 16 September 2026. Public query: <https://dune.com/queries/8750094/12777052/>

A-5. Evidence-record structure

Each illustrative example should be capable of being recorded using the same fields: system or function; source and observation date; production or test environment; workload and topology where known; measured metric; treatment of completion, failure and retry; finality criterion; material dependencies; observed limitation or trade-off; and the narrow prudential inference supported by the evidence. The exhibit above applies that discipline to the extent supported by the underlying public dataset.

Appendix B. Complete Systemic Stablecoin Transaction and Redemption Lifecycle

This is an illustrative lifecycle, not a universal ordering. Stages may coincide, be omitted, recur or occur in a different sequence depending on the architecture and the financial function. The purpose is to prevent a material boundary from being hidden inside a single label such as confirmed, final or redeemed.

1. User intent created
2. Wallet or custodian approval initiated
3. Policy/MPC approval completed
4. Transaction constructed
5. Transaction signed
6. Transaction broadcast
7. RPC or access layer receives transaction
8. Mempool or sequencer receives transaction
9. Ordering process begins
10. Transaction included
11. Transaction executed
12. State transition produced
13. Transaction confirmed
14. Transaction finalised
15. Additional bridge, cross-domain or underlying-layer assurance condition completes, if applicable
16. Issuer redemption system recognises receipt
17. Custodian or issuer books update
18. Bank cash settlement occurs
19. Reserve-side liquidation/settlement occurs, if required
20. Regulator-facing record produced
21. Post-event reconciliation and audit completed

In particular, issuer recognition, book updates, bank settlement, reserve-side activity, regulator-facing evidence and reconciliation need not occur in the numbered order shown. Their actual ordering and dependencies should be mapped for the arrangement being assessed.

The lifecycle should not be collapsed into a single label such as 'instant.' Each material stage can have its own control point, timestamp, dependency, failure mode and potential economic exposure.

The complete path can also include gateways or adapters, orchestration and routing, data-availability or proof dependencies, observers or indexers, cross-domain relay and verification, issuer recognition, redemption orchestration, payment-system settlement and token/cash reconciliation. These stages should be included wherever they are material to the regulated outcome rather than treated as peripheral technical services.

For each material boundary, the arrangement should identify the proposition established, the authority or source for that proposition, how the next component observes it, whether it can change or remain indeterminate, what dependency remains and what action may safely follow.

Appendix C. Measurement Provenance Framework

- What exact component, stage or full lifecycle was measured?
- Was the benchmark measuring authentication, signature verification, propagation, sequencing, execution, proof generation, database persistence or finalised state transfer?
- Were wallet, custody, API/RPC, sequencer, consensus, data availability, bridge, issuer recognition and redemption workflows included?
- Were failed transactions counted as throughput?
- Were retries counted as new transactions?
- Was load realistic, sustained, geographically distributed and production-like?
- What finality criterion was used, and were median and stressed-tail latencies reported?
- Was data availability sized consistently with execution capacity?
- Can the test be independently rerun or corroborated?
- What inference is being drawn, and does it exceed the lifecycle actually measured?

Appendix D. Fit-for-Purpose Assurance Framework

The following framework illustrates one way to translate technology-neutral prudential outcomes into architecture-aware engineering evidence. It is not intended to prescribe a particular distributed architecture, imply that every question applies identically to every arrangement, or substitute for legal and supervisory judgement. Its purpose is to make the suitability decision explicit: what financial function is proposed, what technical arrangement must support it, what properties are required, what evidence demonstrates those properties, and what conclusion follows.

D-1. Define the regulated financial function

Identify the actual workflow and the consequences of delay, ambiguity or failure. Relevant requirements can include normal, peak and stressed workload; completion time; technical and settlement finality; ordering and inclusion; availability; recovery objectives; authoritative-state requirements; redemption or settlement dependencies; and the consequence of missed service levels.

D-2. Map the complete architecture and commit path

Identify the components through which the financial state is established, observed, changed or acted upon, including custody and key management; gateways, adapters and access/RPC infrastructure; orchestration and routing; application or smart-contract logic; the base ledger; Layer 2 or extension layer where applicable; sequencer or ordering service; data-availability or proof layer; bridge or interoperability mechanism; oracle or external-data dependency; observers, indexers or attestation services; issuer-recognition and redemption systems; settlement banks and payment rails; reserve/liquidity providers; and other material off-chain components. Distinguish components that are merely informational from those required to produce, observe, authorize or act upon a valid financial state transition, and identify the bounded proposition for which each component is authoritative.

D-3. Identify operative state, control and accountability

Determine where authoritative operational state is established; who operates nodes or validators; who controls admission, sequencing, upgrades, emergency actions and privileged keys; which entities or affiliates can exercise decisive influence; and what accountability exists when a required component fails or produces an unexpected state.

D-4. Test the five assurance domains

Assess: (1) state-transition integrity, finality and authoritative state; (2) ordering integrity and inclusion; (3) concentration and control authority; (4) continuity, capacity, recovery and migration; and (5) evidence, verification and disclosure. The specific questions and tests should follow the architecture and the financial function rather than a generic blockchain checklist. Where those domains are exposed to deliberate or malicious action, the same framework provides a structural/adversarial-resilience lens: which properties remain true, under what assumptions, and with what evidence.

Where the state-transition system is material to the prudential conclusion, the assessment should also ask whether its behaviour is sufficiently specified to permit meaningful testing or independent verification. This can include the relevant initial or reference state, admissible and prohibited transitions, safety properties, liveness assumptions, fault and timing assumptions, and expected behaviour under message loss, delay, reordering, partition, equivocation, censorship, privileged intervention and recovery. The depth of specification and verification should remain proportionate to the function and systemic importance; the objective is not to prescribe formal methods, but to ensure that critical guarantees can be examined rather than merely asserted.

The Bank need not perform protocol-level engineering tests itself. Depending on materiality and the proposition being relied upon, the evidentiary basis may be produced by the issuer, infrastructure provider, independent technical reviewer or other competent assurance source, provided the scope, assumptions, results and limitations are sufficiently transparent for the supervisory conclusion being drawn.

D-5. Test representative and stressed conditions

Evaluate the complete system under conditions representative of intended operation, including sustained and peak workload, contention, state growth, relevant geography and topology, component or node failure, congestion, recovery, upgrades and material dependency failures. Where a property has only been inferred from a component benchmark, controlled demonstration or theoretical result, that limitation should remain visible.

Where safety and liveness can conflict, testing should identify the boundary explicitly. The assessment should determine whether the system stops before violating the relevant safety property, what condition causes the halt, how unresolved financial instructions are represented, what degradation occurs before and during the halt, and whether restart and state reconciliation preserve the relied-upon guarantees.

D-6. Identify gaps, mitigants and residual dependencies

For each required property, distinguish what is demonstrated, what is supported only by assertion or incomplete evidence, what remains unknown, and what requires further testing. Where a mitigant is proposed, determine whether it actually removes the dependency or merely relocates it to another component, operator, contractual mechanism or recovery process.

D-7. Reach an explicit suitability conclusion

The assessment should permit more than a binary pass/fail result. Illustrative conclusions include: Fit for Purpose; Fit Subject to Defined Conditions or Controls; Production Fitness Not Yet Demonstrated; or Not Fit for the Proposed Function. 'Not yet demonstrated' is distinct from a finding that an architecture is incapable of meeting the requirement; it identifies an evidentiary gap that should be resolved before reliance increases.

D-8. Architecture-aware diligence questions

The same prudential outcome can require different evidence depending on the architecture. The following questions are illustrative and should be applied only where relevant.

Base-layer / L1 arrangements

Who participates in consensus or validation? What fault assumptions and finality semantics apply? How does capacity change under sustained load or congestion? Who can change protocol rules or exercise emergency authority? What production evidence demonstrates the required behaviour?

Layer 2 / extension-layer arrangements

Who sequences or orders transactions? What happens if the sequencer or prover is unavailable? Where is transaction data available? What dependency remains on the base layer? When does state become technically final and financially usable? What withdrawal, escape or recovery mechanisms exist, and who controls them?

Bridges and interoperability mechanisms

What event or proof causes state to be recognised across systems? Who attests to or validates that state? What additional trust, governance or operational dependency is introduced? How are inconsistent states, freezes, delayed messages and recovery handled? Which system is authoritative while the mechanism is impaired?

Permissioned or institutional networks

Who admits and removes participants? Who operates validating nodes and critical shared services? Are there common cloud, software, identity or operator dependencies? Who controls upgrades and emergency changes? How are disputes, partitions, member failure and coordinated recovery handled?

D-9. Evidence status

To avoid converting absence of evidence into a positive assurance conclusion, material findings can be classified as: Demonstrated; Demonstrated Subject to Defined Conditions; Partially Demonstrated; Not Independently Demonstrated; Evidence Unavailable or Incomplete; Requires Additional Testing; or Material Deficiency.

D-10. Supervisory use

The framework is intended to support structured dialogue among issuers, infrastructure providers, regulated institutions, independent reviewers and supervisors. It does not require the Bank to prescribe technical design. It provides a way to ask whether the evidence available for a particular arrangement is sufficient for the financial function and level of reliance proposed.

Appendix E. Systemic Stress Scenarios and Evidence Requirements

Purpose. TauriPay suggests these scenarios as architecture-aware inputs to the Bank's existing extreme-but-plausible testing approach. They are not predictions that a particular technology will fail. They test whether an arrangement can preserve or reconstruct authoritative financial state when components remain partially functional, observations diverge, assurance is delayed, automated workflow decisions occur before an earlier outcome is known, or a material authority or dependency behaves adversarially.

- **Capacity snowball:** adoption increases load, failures and retries amplify demand.
- **Freeze-state mismatch:** orchestration, issuer books and chain state disagree at intervention time.
- **Sequencer or key event:** restart, replay, censorship or prioritisation authority becomes material.
- **Affiliated-operator concentration:** nominal decentralisation differs from decisive control.
- **Ordering scarcity:** queue or fee priority determines who exits first.
- **Challenge-period or layered-finality domino:** provisional credit becomes cross-institution exposure.
- **Release-management failure:** emergency patch changes transaction or economic behaviour.
- **Indexer/RPC divergence:** network is nominally up while institutions observe different state.
- **Redemption queue dispute:** timing and ordering determine which holders enter the redemption process first.
- **Adversarial control event:** compromise or misuse of a sequencer, validator set, administrative key, governance authority or other privileged control tests containment, invalid-state prevention, affected-instruction identification and recovery.
- **Selective dependency disruption:** transaction access, observation or cross-domain communication is selectively impaired while portions of the underlying infrastructure remain nominally available, testing whether delay, censorship, failure and indeterminate outcome can be distinguished.

For each scenario, the supervisor should identify the required evidence packet: transaction state, ordering, finality, failures/retries, control actions, intermediary exposure, activated runbooks and point-in-time reconstruction.

Appendix F. Consultation Question Crosswalk

Q1: Current Code effect: Sections I, III, VI, VIII, X and XIII.

Q2: Financial-stability consequences: Sections I, II, III, V, VI, VII, IX, X, XII and XIII.

Q3: Operational constraints: Sections II, III, IV, V, VI, VII, VIII, IX, XI and XIII.

Q4: Other comments / developing framework: all analytical sections and Appendices A-H.

Q5: Equality Act 2010: TauriPay has not identified a specific differential impact on persons sharing protected characteristics arising from the infrastructure-evidence proposals in this response. No separate equality-impact recommendation is made.

This crosswalk is intentionally high-level, but the substantive allocation is important. Q1 supports limited current-Code clarification around redeemability, dependencies, recovery and material change. Q2 supports the financial-stability rationale for examining how infrastructure stress can propagate into redemption, liquidity and confidence effects. Q3 is the principal home for architecture-aware operational constraints and extreme-but-plausible scenario design. Q4 provides the principal route for recommendations that are better developed through future supervisory materials rather than embedded immediately as detailed Code rules. Q5 is addressed separately above.

Appendix G. Prudential Standards Crosswalk

Purpose. This crosswalk shows how the response translates established prudential outcomes into observable operational evidence. It does not propose a separate supervisory regime or substitute technical criteria for the Bank's statutory and supervisory judgement.

G-1. Governance and accountability. The Bank's June 2026 policy and draft Code place responsibility on the systemic issuer within a jointly regulated framework. CPMI-IOSCO applies PFMI governance and comprehensive-risk-management expectations to systemically important stablecoin arrangements, while the FSB calls for clear governance and accountability across relevant functions and activities. The corresponding evidence in this response is a functional control and affiliation map covering decisive actors, privileged authority, material service providers, escalation and responsibility for recovery.

G-2. Settlement finality and authoritative state. CPMI-IOSCO identifies settlement finality and money settlements as core PFMI considerations for systemically important stablecoin arrangements. The Bank has also stated that its developing regime will include application of the PFMIs to systemic stablecoin issuers. The proposed evidence therefore distinguishes execution, confirmation, technical finality, issuer recognition, settlement finality and legal consequence, and asks which record or determinable state is authoritative for the relevant obligation. The September 2026 FCA/Bank Feedback Statement also records respondent concern that contractual finality may not provide statutory insolvency protection and recognises the importance of settlement and wider legal certainty.

G-3. Redemption and practical performance of the monetary claim. The Bank's regime gives direct and timely redemption a central role, and the FSB likewise emphasizes timely redemption at par. This response adds the infrastructure question: whether the complete redemption path can carry holder instructions through the material stages required for issuer receipt and recognition under representative and stressed conditions. Redemption-effective capacity is the proposed function-specific measurement.

G-4. Operational resilience, recovery and continuity. PFMI operational-risk principles require robust systems, capacity, continuity and recovery. The Bank's operational-resilience framework already requires end-to-end mapping, interdependency analysis, impact tolerances and extreme-but-plausible scenario testing. This response does not duplicate those requirements. It adds an architecture-aware state and assurance lens for scenarios in which components may remain available while the financial obligation becomes partial, differently observed or indeterminate, requiring recovery to an authoritative financial state rather than technical restart alone.

G-5. Capacity, scalability and measurement provenance. The Bank's DLT Innovation Challenge expressly treats scalability, settlement finality, control and interoperability as interdependent design questions and cautions that illustrative challenge results were not independently validated. This response operationalizes the same discipline by separating claimed or theoretical capacity, observed gross activity, successful activity, assured or finalised activity, complete regulated outcomes and degraded-state capacity. The conclusion drawn from a test should not exceed the lifecycle and conditions actually measured.

G-6. Interdependencies and third-party risk. The PFMI, FSB and Bank frameworks already require treatment of material dependencies and third-party risks. The additional proposition here is boundary-specific: for a material dependency, identify not only that the component exists and can disrupt service, but what financial proposition it establishes, whether it is authoritative for that proposition, how the next component observes it, what remains unresolved and what action can safely follow. The proposed Infrastructure Dependency Assessment applies that discipline across ledgers, sequencers, bridges, custody, key management, access layers, cloud or software dependencies and other material services. The same statement identifies bridges, key-management systems, governance arrangements, oracle providers and cross-chain messaging layers as areas of operational-resilience concern.

G-7. Continuing evidence and material change. The Bank has indicated that supporting materials on reporting, record-keeping, operational resilience and the application of the PFMIs will continue to develop. The response therefore treats systemic suitability as a maintained conclusion rather than a one-time approval event. Production metrics, incident evidence, governance actions, recovery exercises and material changes to consensus, sequencing, bridges, software, keys, dependencies, capacity or recovery mechanisms should remain capable of reassessment where they affect a relied-upon property.

G-8. Proportionality and technology neutrality. The Bank's policy uses a proportionate, risk-based approach and invokes the principle of the same risk, same regulatory outcome. Consistently, this response does not rank public, private, permissioned, permissionless, layered or conventional architectures. The supervisory inquiry follows the financial function and the material dependency. Different architectures may satisfy the same prudential outcome in different ways, but where their operating characteristics differ materially, the evidence supporting reliance should remain visible. It further states that equivalent treatment of tokenised and non-tokenised assets depends on the risks being comparable and sufficiently mitigated.

Primary standards and policy sources. Bank of England, Sterling-denominated systemic stablecoins: Policy Statement and consultation on draft Code of Practice, 22 June 2026, <https://www.bankofengland.co.uk/paper/2026/ps/sterling-denominated-systemic-stablecoin>; CPMI-IOSCO, Application of the Principles for Financial Market Infrastructures to stablecoin arrangements, 13 July 2022, <https://www.bis.org/cpmi/publ/d206.htm>; Financial Stability Board, High-level Recommendations for the Regulation, Supervision and Oversight of Global Stablecoin Arrangements, 17 July 2023, <https://www.fsb.org/2023/07/high-level-recommendations-for-the-regulation-supervision-and-oversight-of-global-stablecoin-arrangements-final-report/>; Bank of England and BIS Innovation Hub London Centre, DLT Innovation Challenge 2025: Final Report, 12 May 2026, <https://www.bankofengland.co.uk/research/fintech/dlt-innovation-challenge-2025>. Financial Conduct Authority and Bank of England, FS26/1: Tokenisation in Wholesale Financial Markets, 14 September 2026, <https://www.fca.org.uk/publication/feedback/fs26-1.pdf>.

Appendix H. Questions Before Systemic Reliance

The questions below provide a front-end diagnostic for identifying what is known, what has been demonstrated and what remains unresolved before reliance on infrastructure reaches, or increases to, a systemically material level. They are illustrative rather than prescriptive and do not prescribe a particular technology, architecture or answer. Not every question will apply to every arrangement, and different technical designs may satisfy the same prudential outcome through different mechanisms. They are not a proposed mandatory checklist or separate technology regime, and do not require the Bank to certify an underlying protocol. The accountable institution should instead be able to demonstrate a proportionate evidentiary basis for material reliance.

For each question, distinguish what is demonstrated by current production evidence from what is supported by controlled testing or simulation, projection or assumption, assertion, or information that is not presently available. Where a function is performed or represented by a service provider or other intermediary, the question applies to the underlying function and state on which that service relies.

A. Demonstrated performance

1. What level of transaction activity has the system actually sustained in production, for how long, and with what rates of successful completion, failure and re-attempt?
2. How does that demonstrated performance compare with the capacity commonly attributed to the system?
3. From what event to what event is performance actually being measured, and what stages or dependencies fall outside that measurement?
4. As demand increases, at what point do completion times, failures, re-attempts or unresolved transactions begin to increase materially?

B. State, time and finality

5. When a transaction is first acknowledged, what has actually occurred, and what material steps remain?
6. When the system reports that a transaction is confirmed or final, what exact state or event has been established?
7. Can that state still change? If so, what could change it, for how long, and what determines when that possibility ends or becomes sufficiently remote?
8. After a transaction is reported as complete or final, what events, conditions or dependencies, if any, remain before other participants can safely rely on that state?
9. What establishes confidence that the state will not change, and does the system ever reach a point at which it can no longer change under its operating rules?

C. Authority, ordering and control

10. How many independently controlled participants actually perform the function that establishes the state being relied upon, and what determines who is able to perform that function?
11. Who or what determines transaction order, and can that function delay, exclude or report acceptance before the state ultimately relied upon has been established?
12. Can a smaller group exercise consequential powers that the broader network cannot, including changing operating rules, pausing processing, upgrading the system or exercising emergency authority?
13. Can the ability to observe, insert, reorder, delay or exclude transactions materially affect financial outcomes, and who, if anyone, possesses that capability?

D. Represented and authoritative state

14. When an institution or its service provider reports that a transaction has succeeded, completed or become final, what underlying state is that report based upon, and is that state authoritative for the proposition being relied upon?

E. Failure, uncertainty and recovery

15. During a disruption, can participants determine which transactions reached the relevant state, which failed and which remain unresolved, and what action is safe for each?
16. If the system failed after a state had become authoritative, could that same state be recovered or reconstructed, including changes occurring after the latest recoverable record?

An unknown or incomplete answer does not by itself establish that an infrastructure is unsuitable. It identifies a matter whose significance, evidentiary basis and any necessary mitigation should be understood before reliance on that infrastructure reaches, or increases beyond, a level supported by the available evidence.

Detailed Assessment Questions

The questions above provide a front-end diagnostic. Where an answer is unknown, incomplete, qualified or insufficiently evidenced, the more detailed questions below illustrate areas of inquiry that may assist in understanding the relevant architecture, dependency, state, authority, performance and recovery characteristics. The depth of inquiry should remain proportionate to the extent to which a component or dependency can alter, delay, obscure or prevent establishment of the financial state required for the regulated outcome.

H-1. Ledger placement and authoritative state

- What ledger, domain or combination of domains establishes the state on which the relevant financial function relies?
- Does relevant activity occur on an L1, L2, sidechain, application chain, permissioned ledger or other execution domain?
- For each material proposition, which layer establishes authoritative state and which layers merely execute, observe, relay or represent it?
- Can the instrument or a claim relating to it be represented on more than one domain? If so, is one representation canonical?
- Can different domains or intermediaries temporarily present different states for the same financial proposition?
- What event permits an institution to treat the relevant state as sufficiently established for the next financial action?
- Can an intermediary report completion before the state that is authoritative for the relevant proposition has been established? If so, what proposition is the intermediary representing and how is subsequent divergence detected and reconciled?

H-2. Consensus, liveness and fault-threshold evidence

- What consensus or state-establishment mechanism is relied upon, and what safety and liveness thresholds are claimed?
- What behaviour is expected as participation or connectivity degrades toward the claimed fault threshold?
- Has the arrangement been tested under conditions approaching, reaching and, where meaningful, exceeding material claimed safety or liveness thresholds or operating boundaries relevant to the architecture?
- Do tests include correlated failures, partitions, delayed messages, validator concentration or other conditions that may differ from independent-node failure assumptions?
- At each tested threshold, does the system continue, degrade, halt, fork, delay finality or enter another bounded state?
- What evidence supports these conclusions: specification, formal analysis, simulation, controlled adversarial testing, testnet evidence, production observation or some combination?
- If safety is preserved while liveness is lost, how is that condition detected and what financial state can participants safely rely upon while the system is stalled?

H-3. Sequencing, ordering and inclusion

- Is transaction ordering performed by a sequencer, ordering service, leader or equivalent mechanism? Who operates or controls it?
- What does acceptance by the ordering mechanism establish, and what does it not establish?
- Where applicable, how are sequencer acceptance, execution, L1 publication or inclusion, assurance/finality and financial completion distinguished?
- Can the ordering mechanism delay, censor or reorder transactions, and can such behaviour materially alter the financial outcome?
- What happens to transactions accepted immediately before sequencer or ordering failure?
- Is there a bypass, fallback or forced-inclusion mechanism? Under what conditions is it available and has it been tested?
- How are pending, rejected and indeterminate transactions represented during an ordering interruption?

H-4. Smart contracts, privileged control and governance

- Which contracts or programmable components establish or materially affect the state of the instrument?
- Are those components immutable, upgradeable, pausable, replaceable or subject to emergency intervention?
- Who controls administrative, upgrade, pause, mint, burn, freeze or other privileged authorities, and what information is available to the issuer and, where relevant to supervision, to the Bank about the persons or entities able to exercise those authorities?
- What authorization threshold applies to privileged actions, including multisignature or committee arrangements?
- Are privileged changes subject to timelocks, and can those timelocks be bypassed in an emergency?
- What happens if a privileged authority, administrator or key is compromised or unavailable, and what containment and recovery mechanism prevents or bounds unauthorised state change?
- What audit trail establishes who exercised privileged authority, when it was exercised and what state changed?
- What happens to pending transactions and dependent representations during upgrades, pauses or emergency interventions?

H-5. L2, bridge and cross-domain state

- For each cross-domain boundary, what proposition is established on the source domain and what assurance is required before another component acts upon it?
- What observes, relays or proves source state, and what verifies that representation on the destination domain?
- What establishes destination execution, and what additional conditions remain before end-to-end completion?
- Can source and destination domains disagree or remain at different assurance statuses or thresholds for a material period?
- What happens if a bridge, relay or proof mechanism stops after source action but before destination action?
- Can the same financial proposition exist simultaneously in ambiguous, duplicated or otherwise inconsistent representations?
- Which state is authoritative during disagreement, and how is that determination communicated to dependent institutions?
- How are source, destination and intermediary records reconciled following recovery?

H-6. Oracles, RPCs, indexers, observers and represented state

- Does the financial workflow depend on an oracle, RPC provider, indexer, explorer, observer or other derived representation of ledger state?
- For what bounded proposition, if any, is each source authoritative?
- What is the observed latency between authoritative state and the representation consumed by the institution?
- For a material representation, what are the underlying event time, observation time, report time and reliance or action time?
- How current must the representation be, and what assurance threshold must it satisfy when acted upon? Can a representation that was accurate when observed become stale or insufficient before the next financial action?
- Can observations become stale, incomplete, concentrated or unavailable? Can different observers disagree?
- When observations conflict, what hierarchy or reconciliation process determines the state on which the institution may safely act?
- Could an institution take financial action based on a representation that later proves inconsistent with the state authoritative for the relevant proposition?
- Is provenance retained so that the institution can establish what information it relied upon at the time of a material action?

H-7. Orchestration, correlation and retry control

- Is an orchestration layer used to route, correlate, time out, retry, fail over or otherwise manage distributed-infrastructure activity?
- What state does the orchestration layer maintain, and which external propositions does it consume?
- Which consumed propositions are authoritative and which are observations or intermediate acknowledgements?
- When an intermediary reports a downstream state, is it asserting a proposition within its own authority, representing a proposition established by another system, or contractually standing behind that representation; and is that distinction visible to the institution relying upon it?
- How does orchestration distinguish timeout from known failure and from an indeterminate destination outcome?
- What prevents a retry or resubmission from creating duplication where the original outcome remains unresolved?
- How is orchestration state reconciled with authoritative destination state following disruption, failover or delayed observation?
- Can the orchestration layer itself become a source of divergence between institutional records and chain state?

H-8. Performance, capacity and measurement provenance

- What throughput or transactions-per-second figure is claimed for each material layer, and what exactly does the figure count?
- Does the measure represent submitted, accepted, ordered, executed, successful, failed, retried, included, finalized or otherwise assured activity?
- Are protocol votes, system transactions, failed attempts or retries included in the headline figure?
- Is the cited result a peak, burst, average or sustained measurement, and over what duration and transaction complexity?
- What successful throughput has been observed at the assurance level required for the financial function?
- What happens to successful capacity, latency and completion probability during congestion and, where relevant, as the architecture approaches material operating or fault thresholds?
- How many infrastructure actions are required to complete one financial obligation, and can retries or cross-layer activity amplify load under stress?
- Can uncertainty generate additional query, observation, reconciliation or provider-failover demand, and has the observation path been tested under the same stress conditions that increase that demand?
- Which external dependencies constrain otherwise available ledger capacity?
- Can the measurement be reproduced, and are its definitions, denominator, workload, topology, software version and observation period retained?

H-9. Failure, snapshotting, state reconstruction and reconciliation

- What state, if any, is persisted independently of live infrastructure for recovery purposes, including through snapshots, checkpoints, journals, logs or equivalent recovery records?
- What point in the transaction or state-transition lifecycle does each recovery record represent?
- What state is captured, at what interval, by whom, and can the recovery record be verified against a source authoritative for the proposition it is intended to preserve?
- What state changes may occur after that recovery point but before the next durable recovery record is established, and where are those intervening changes preserved?
- What is the maximum interval of state that may require reconstruction after disruption?
- How is it determined which post-recovery-point transactions or state transitions had become authoritative for the relevant proposition before disruption?
- Can intervening state transitions be replayed, rederived or otherwise reconstructed in a manner consistent with the architecture?
- Can the relevant authoritative state be reconstructed without relying solely on the failed component?
- Could a participant or intermediary have been permitted to rely upon a transfer as completed before the state necessary to reconstruct the relied-upon proposition was durably recoverable?
- How are transactions occurring around the recovery boundary treated, including pending, rejected, failed and indeterminate outcomes?
- How are duplicate attempts, retries or resubmissions identified during reconciliation?
- How are institutional records reconciled against recovered ledger or system state, and where relevant how are L1, L2, bridge, oracle and orchestration representations reconciled with one another?
- If restored infrastructure state conflicts with an intermediary's or institution's pre-disruption representation, which source is authoritative for the proposition at issue and how is the discrepancy resolved?
- What recovery-point and recovery-time objectives apply specifically to restoration of the relied-upon financial state and reconciliation, rather than merely restart of infrastructure?
- Has state reconstruction and reconciliation been tested under realistic failure conditions, or is the process documented but untested?

H-10. Continuing assurance and material change

- What architecture, contracts, consensus parameters, sequencer configuration, bridge, oracle, observation path, node topology and recovery mechanisms were within the scope of the evidence supporting the current suitability conclusion?
- What material assumptions concerning capacity, fault and adversarial tolerance, finality, control integrity or dependency concentration underlie that conclusion?
- What changes require reassessment, including contract upgrades, new execution layers, bridge replacement, sequencer changes, consensus changes, validator concentration, oracle changes, material workload increases or new cross-domain representations?
- What governance, change-control and reporting process identifies a material change in a dependency or authority boundary, and how would that change be surfaced to the accountable institution and, where applicable, the supervisor?
- Does continuing evidence remain sufficient for the scale and systemic reliance actually reached, rather than the scale originally assessed?

A systemic suitability conclusion should not exceed what the available evidence demonstrates. The questions above are intended to make that discipline operational at the points where distributed state, dependent representations and cross-layer assurance become material to the regulated financial outcome.