AUGUST 2026

APPLIED QUANTUM

CRYPTOGRAPHIC CONCENTRATION IN PAYMENTS

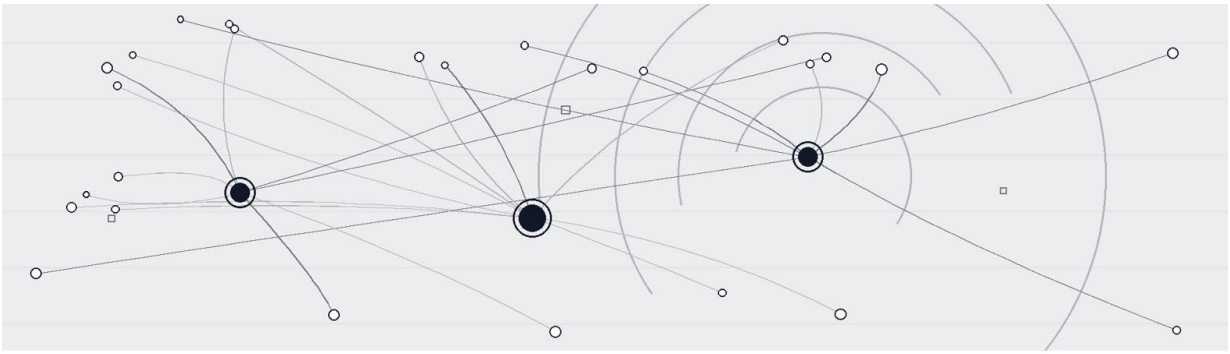

Measuring cryptographic concentration beneath the vendor layer

Version 1.0-RC – August 2026

Steve Vaile and Marin Ivezic, Applied Quantum

"Switching vendors moves the contract, not the dependency."

COPYRIGHT AND LICENSE

DISCLAIMER

This document is provided as professional guidance based on the authors' and Applied Quantum's practitioner experience and publicly available standards, regulations and industry research. It does not constitute legal, regulatory, financial or compliance advice. Organisations should consult qualified legal counsel, auditors and regulatory authorities for guidance specific to their jurisdiction, sector and circumstances.

The regulatory instruments, certification records, standards and vendor capabilities referenced in this document reflect the state of the landscape as of August 2026, and that landscape moves. Readers should verify current status against primary sources (regulators and supervisory authorities, NIST, EMVCo and PCI SSC, the CycloneDX specification, public certification listings, vendor documentation) before making decisions.

References to specific vendors, products or tools are for illustrative purposes and do not constitute endorsement.

This is a v1.0 release candidate: v1.0 final publishes in November 2026 after a pilot cycle. Open items are public on the roadmap at CCFramework.org, and corrections are published as dated entries on its errata page.

ABOUT THIS DOCUMENT

Document type	Practitioner whitepaper
Parent document	Companion to The Cryptographic Concentration Framework – Universal – v1.0-RC (August 2026)
Intended audience	Chief risk officers, heads of operational and third-party risk, payment-scheme and supervisory readers, and the practitioners who brief them
Assumed knowledge	Payments-industry context; no framework knowledge is assumed
Scope	The argument and the evidence: why vendor-substitutability tests cannot detect shared-lineage concentration, what the public record shows, and what a payment institution can measure this quarter. The method is summarised here and defined normatively in the Universal

	Framework.
Status	Version 1.0-RC – release candidate; v1.0 final November 2026 after pilot

VERSION HISTORY

Version	Date	Changes
1.0-RC	August 2026	First public release, as a release candidate; drafted through four external review rounds and a family-wide line-edit pass.
1.0-RC	20 August 2026	Co-author review edits accepted (Steve Vaile, 19–20 August).
1.0-RC	7 September 2026	Three citations moved to primary sources; Section 10 pointer to the Payments Extension restated.
1.0-RC	8 September 2026	Retitled Cryptographic Concentration in Payments, subtitle “Measuring cryptographic concentration beneath the vendor layer”. Section 1 dates the first four UK Critical Third Parties designations (in force 13 July 2026); the Project Leap verification ratio reads about 7.5 times; the NIST IR 8547 sentence names the 112-bit strengths deprecated after 2030 and the report’s draft status; the worked-example table reports the largest failure domain per layer. Prose revised throughout. No figure changed.

HOW TO USE THIS DOCUMENT

This paper stands alone. It makes the argument, shows selected worked figures and states what a payment institution can do this quarter. The measurement method it summarises is defined normatively in the CCF Universal Framework, and the payments enumerations behind the worked example live in the Payments Extension.

ACCOMPANYING RESOURCES

The framework family publishes at CCFramework.org: the Universal Framework, the Payments and Financial Services extensions, the payments whitepaper, the CBOM Conformance Statement, the Technical Companion and the instruments, together with the public roadmap, the errata and corrections page, the provenance survey and the frontier-census methodology.

Every canonical figure in the family comes from the reference implementation, with ten canonical test vectors that fix each one exactly. Both are open under Apache-2.0 on the Applied Quantum GitHub, and any implementation that reproduces the fixtures is conformant.

The Applied Quantum CBOM Profile, the input data contract, publishes at CBOMProfile.org. The Applied Quantum PQC Migration Framework, whose sector taxonomy the extensions follow, publishes at PQCFramework.org. Practitioner background and adjacent analysis sits on PostQuantum.com, as further reading rather than a normative source.

TABLE OF CONTENTS

Disclosure. Applied Quantum sells post-quantum migration advisory services. This paper concludes that cryptographic concentration is a systemic exposure requiring an accountable owner, which favours our commercial interest. Three things are offered against that. The framework, its assessment method, and its instruments are published openly under Creative Commons Attribution 4.0. An institution can run the assessment with its own staff, on its own inventory, without engaging us. And every claim below cites a public primary source inline, not proprietary modelling.

SUMMARY

A payment institution can operate four hardware security module vendors, three TLS terminators and two certificate providers and pass every supplier-level substitutability test. Resolving those relationships to the components executing underneath can substantially reduce the apparent diversity. Almost all of the suppliers implement RSA and elliptic-curve cryptography. At the algorithm layer, that estate is one dependency, and adding a fifth HSM vendor does not change it.

The **Cryptographic Concentration Framework (CCF)** measures these dependencies at the level of an important business service. It records concentration across six layers, the reach of shared failure domains, and the operational consequences against the service's tolerances. **Effective Coverage** measures whether every relevant party and hop operates at the defined cryptographic target state. Together, these outputs distinguish a procurement issue from an implementation, recovery or counterparty-coordination issue.

Four arguments support the assessment. First, DORA Article 29 and the UK Critical Third Parties tests assess substitutability at the vendor layer. They cannot identify the cryptographic commonality beneath those suppliers through that test alone.

Second, the incident evidence is documented. KyberSlash, disclosed from late 2023 to early 2024, affected twelve libraries. In the cases traced by the researchers, vulnerable divisions travelled by inheritance from one reference implementation; independently written decapsulation paths were unaffected. Shared implementation lineage is therefore an observable failure mechanism, not only a concern about a future quantum event.

Third, inventory practice does not ask the necessary questions. For lineage, the standard has the field but discovery tools do not populate it. For key-generation design, the required field does not exist in the core model. A standards-conformant inventory can consequently omit the evidence needed to resolve common failure domains.

Fourth, the migration deadline can create the correlation it is intended to reduce. Institutions move towards one horizon using a small number of immature libraries under

the same time pressure. Algorithm migration and implementation concentration therefore need to be assessed together.

The framework is not a security rating for algorithms or vendors. It does not estimate the probability of a cryptographic break, translate concentration into financial loss, or certify compliance with a regulatory regime. It provides a measurement method a second-line function can run this quarter and a basis for assigning one accountable owner to the resulting exposure. The worked example is synthetic, not a market benchmark or a measured institution.

1. SUPPLIER RELATIONSHIPS AND SERVICE DEPENDENCIES

Diversification is the first defence in finance, but it depends on separating the relevant failure mechanisms. Institutions can route through a small set of providers that all implement RSA and elliptic-curve cryptography. A break at the algorithm layer then affects every provider at once. The weakness in diversification is structural: the supplier relationships differ while the cryptographic assumption remains common.

A payment institution usually encounters cryptography through operational responsibilities: protecting a PIN-processing environment, validating a certificate chain, authorising a settlement instruction, establishing an encrypted channel, or maintaining keys in a hardware module. Those responsibilities are distributed across teams and suppliers. The resulting inventories tend to reflect the management boundary of the team that maintains them.

A custody inventory identifies modules, partitions and keys. A certificate inventory identifies issued certificates and expiry dates. An application inventory identifies software components. A supplier register identifies contractual relationships. Each can be accurate within its own scope while leaving the institution unable to answer a service-level question: which functions would be affected if one shared cryptographic component failed?

CCF starts with the service because that is where the institution can describe the consequence. Card authorisation, payment release and settlement access do not necessarily have the same tolerances, fallback arrangements or exposure to irreversible outcomes. An implementation used in all three may create a common failure domain, but its operational effect must be established separately for each service. Counting the implementation once in an enterprise technology inventory does not replace that service mapping.

The service boundary also determines which external dependencies matter. An acquiring processor may operate an authorisation service that depends on scheme connectivity and an acceptance-device population. A broader service boundary may also include settlement submission and therefore involve central-bank infrastructure. The framework

does not prescribe a single organisational service taxonomy. It requires the selected boundary to be explicit, linked to an accountable owner and used consistently when comparing results.

Shared infrastructure is counted in every service it supports. A common HSM family does not become one tenth as important to each service merely because ten services use it. This is a service-exposure convention, not an assertion that the enterprise owns ten copies of every key. The distinction must remain visible in institution-level aggregation: a service-weighted figure is not a de-duplicated enterprise asset count.

Substitutability and regulatory scope

DORA has applied since January 2025. Article 28(4)(c) requires an institution to assess whether an arrangement reinforces concentration risk. Article 29 requires assessment of whether a critical provider is "not easily substitutable" and of reliance on the same or closely connected providers.

The UK's Critical Third Parties regime took effect on 1 January 2025 and reaches a provider only on HM Treasury designation. The first four designations – Amazon Web Services EMEA, Google Cloud EMEA, Microsoft Ireland Operations and Oracle Corporation UK – took effect on 13 July 2026. The regime adds to firms' own third-party duties.

These tests operate at the vendor layer. Before migration, switching HSM vendor changes the commercial relationship while leaving the breakable algorithm in place. During migration, providers share immature post-quantum libraries, implementation pitfalls and incomplete inventories. Supplier substitution does not resolve either form of commonality.

Each regime is written around the failure of one provider at a time. A cryptographic break is correlated and simultaneous. Applied at the algorithm layer, DORA's substitutability test identifies the whole market as one dependency, a result the provider-level templates cannot express. CCF supplies a structured cryptographic dependency view within the broader assessment. It is not a regulatory requirement to use CCF, and neither regime should be read as preventing analysis below the direct supplier.

The practical result is a more precise statement of diversification. An institution should be able to say that two platforms have independent operating arrangements while acknowledging a shared implementation ancestor, or that they use distinct implementation families while their generation provenance remains unresolved. That is more useful than treating independence as a single yes-or-no property of the supplier relationship.

Harm channels

A cryptographic failure can turn a loss of trust into a loss of availability. Once signatures cannot be relied upon, counterparties have a reason to quarantine connections and withdraw from processing. The principal effect in that scenario is stopped payment infrastructure, rather than money stolen directly.

The four harm channels differ in scale. Confidentiality exposure is already open: data copied today can be decrypted later. Authentication exposure arises when signatures can be forged. Availability exposure follows when settlement and liquidity stop because the messages supporting them cannot be trusted. Compliance has a different basis: an obligation falls due on a fixed date whether or not a break occurs. Authentication, the channel against which most payment controls are designed, is the smallest of the four in this analysis.

2. INCIDENT RECORD

The case for examining shared cryptographic dependencies does not depend solely on a future quantum event. The public record includes failures in implementation and key generation that crossed product boundaries. These incidents support the need for dependency analysis, but they do not supply a current census of shared code or establish the vulnerability of today's deployed products.

ROCA and generation provenance

In 2017, a flaw in Infineon's RSA key-generation library, documented in ROCA and public disclosure, affected Infineon trusted platform modules, YubiKey authentication tokens, Gemalto smartcards, BitLocker and the Estonian national identity scheme. Estonia suspended the certificates of 760,000 ID cards in November 2017. The products had their own brands, supply chains, certifications and assessments, but shared one library. Payments is more exposed to this mechanism because secure elements in issued cards come from the same small group of semiconductor vendors.

For CCF, the important distinction is between the key's current location and its generation history. Moving a weak key to a different custody platform does not repair the key. An institution assessing only the platform performing today's operations can therefore miss an inherited exposure in previously generated material. A generation-origin record and the design evidence behind it answer a different question from the current custody register.

ROCA also illustrates why certification and independence must not be conflated. A validation record can help identify a module and its operating context. It does not, by itself, demonstrate that two modules have independent generation designs. The response to a shared generation failure may require replacing an issued population,

redistributing trust and addressing past confidentiality or integrity exposure, not simply installing a firmware update.

The Universal Framework's Annex 1 uses ROCA to examine the layer-6 reading against the public record. Its validation claim concerns the availability of pre-disclosure dependency inputs, not a prediction of the cryptanalytic flaw. Section 12 sets out that distinction.

KyberSlash and shared code

Over December 2023 and January 2024, researchers disclosed KyberSlash: secret-dependent division timing in the decapsulation path of ML-KEM implementations. The academic write-up describes timing vulnerabilities in several implementations, including the official reference code. Kannwischer and colleagues name twelve libraries: zig, kyber-k2so, CIRCL, AWS-LC, Botan, liboqs, crystals-go, PQClean, pqcrypto-kyber, pypqc, pqm4 and kyberlib.

In the traced cases, vulnerable divisions came from the reference implementation by inheritance. Where common ancestry has not been established, the framework's rule still applies: the same bug is not by itself proof of the same lineage. Independently written decapsulation paths, including BoringSSL's, were unaffected. Bernstein's disclosure to the NIST forum reports key recovery against end-2023 reference code on a Raspberry Pi 2. The relevance is implementation behaviour, not a break of the algorithm's mathematical security assumption. Historical affected-code findings remain distinct from the status of a current product version.

A lineage assessment examines whether the relevant code was inherited, incorporated as a dependency, replaced or materially changed. The existence of similar vulnerabilities in two projects is not sufficient to establish a common ancestor. Conversely, different project names and separate maintenance teams do not establish independent implementation where both retain the same upstream routine.

This distinction determines what should be recorded. An implementation may be a fork of another codebase, in which case ancestry is relevant. A product may instead link to a library, in which case a dependency relationship is relevant. Both can create shared exposure, but they are not the same relationship. Treating every dependency as an ancestor produces a graph that looks detailed while misrepresenting how a defect could propagate.

Compiler-induced leakage

In 2024, PQShield's Antoon Purnal reported a related implementation issue. Recent Clang versions recognised the reference implementation's `poly_frommsg` function as a bit test and generated a secret-dependent branch absent from the source. The source was constant-time; the binary was not. Every library using that function verbatim was

exposed. Source-level review and a separate supplier name did not remove the shared toolchain issue.

Post-quantum implementation sources

PQClean, a common ancestor of a large share of deployed post-quantum code, was archived read-only on 4 August 2026 after a January retirement announcement. Its successor also concentrates implementation supply. The PQ Code Package's mlkem-native is the default ML-KEM implementation in liboqs and AWS-LC, and by inheritance in rustls.

Independent implementations exist. OpenSSL 3.5 wrote its own, with a design document addressing constant-time compilation. wolfSSL removed its liboqs integration in favour of native code. BoringSSL, the Go standard library, Bouncy Castle and Microsoft's SymCrypt maintain separate lineages.

The proposition that the number of independent ML-KEM code families in wide financial deployment is "in the single digits" is what the frontier census tests; it is not an established count. The methodology is at ccframework.org/census, and the implementation list is on the provenance page. The supplier disclosure programme corrects that list where it is wrong.

Debian and generated-key exposure

The Debian OpenSSL advisory of May 2008 described predictable random-number generation caused by a Debian-specific change. Its significance for assessment is not limited to the software installed at the time of assessment. Keys generated under an affected implementation can remain in use after the software itself has been corrected or the keys have been moved elsewhere.

A service assessment therefore needs to connect active key material to the relevant generation point and period. Current-version discovery alone may establish that the executing library is fixed while leaving historical keys unresolved. The response can require both software correction and key replacement, with different owners and evidence.

Taken together, these cases support a defined analytical task: identify the implementation, trust and generation relationships through which a common failure could reach a service. They do not establish that every shared component is unsafe or that independent code is inherently preferable. Implementation quality, assurance and operational maintainability remain separate considerations.

3. SIX-LAYER VIEW

CCF separates dependency layers because the unit that matters changes with the failure mechanism. Algorithm names are sufficient for one question and inadequate for another. A vendor name is useful commercial context but rarely resolves the cryptographic failure domain on its own.

Layer	Primary unit used for the concentration index	Additional relationships examined through reach
1. Algorithm	The quantum-vulnerable public-key class under the framework's Shor-type scenario	Reported as a separate migration baseline
2. Implementation lineage	Executing implementation family	Ancestor codebases and shared implementation dependencies
3. Trust root	Terminating anchor of the primary accepted certification path	Alternative paths, shared authorities and policy dependencies, under a stated failure mode
4. Key custody	Relevant firmware family	Shared manufacturer, hardware or software domains where the scenario and evidence justify them
5. Protocol and negotiation	Observed, function-qualified cryptographic outcome	Shared downgrade or negotiation mechanisms where established
6. Key generation	Identified generation point	Entropy-source and key-generation designs

Cryptographic inventory representation

The measurement gap begins with what inventories are designed to record. CycloneDX, published as ECMA-424 and current at version 1.7, is the base model. Its `cryptoProperties` object models algorithm identity and parameters, certificates, protocols and key material. Algorithm, trust root and protocol – three of the six CCF layers – are expressible in that model.

Lineage is an adoption gap. CycloneDX's `pedigree` carries ancestors, descendants and variants and is intended to represent forks. The structure exists, but no discovery tool populates it and inventory practice does not ask for it. Closing this gap does not require a change to the specification.

Key generation is a schema gap. A deterministic random bit generator's presence can be recorded, but the core model does not record the noise source and generation-algorithm design required here. ROCA was a defect in that design. An inventory can therefore be conformant while omitting the information needed to assess this failure domain.

Applied Quantum publishes the CBOM Profile, a vendor property taxonomy layered on CycloneDX. Its namespace registration was filed on 14 August 2026 and is reserved as CycloneDX taxonomy issue #189. The Profile carries implementation ancestry, entropy design, custody firmware family and service mapping. Two additions are proposed for the core specification because vulnerability propagation is a general use case: KyberSlash, Heartbleed and Log4Shell all travelled through ancestry.

Algorithm baseline

For layer 1, CCF groups public-key assets dependent on integer factorisation or discrete logarithms under the specified quantum cryptanalytic failure mode. This is a measurement convention. It does not claim that RSA, finite-field Diffie-Hellman and elliptic-curve systems have identical classical security properties, attack costs or operational uses.

Because the in-scope population is grouped into one class, its CCI is 10,000 while any such assets remain. That number is not useful for distinguishing supplier diversification during migration. CCF therefore reports layer 1 separately and excludes it from the service's structural maximum across layers 2 to 6. The migration information is the size and composition of the remaining baseline population, together with Effective Coverage and the service consequence assessment.

A sound hybrid construction can meet an approved target while retaining a classical component in that baseline. There is no contradiction: the baseline records the component's presence, whereas the target-state determination considers whether the complete construction provides the required protection. Breaking one component of a sound hybrid does not automatically break the protected function.

Executing identity and shared ancestry

Layer 2 first identifies the implementation family actually performing the function. It then records shared upstream relationships separately. Two different executing families remain different units in the index even when a common ancestor reaches both. The reach calculation makes that overlap visible without forcing a non-exclusive dependency graph into an exclusive concentration formula.

Evidence of material divergence can change the relevant ancestry relationship. The fact that a fork is old, separately branded or independently maintained is not sufficient. The assessment needs evidence concerning the code and primitives relevant to the failure under consideration. A project-wide genealogy is useful context, but it is not a substitute for identifying the deployed implementation and applicable code path.

Trust roots and accepted paths

At layer 3, an issuer name is not a complete dependency identity. The assessment must establish the trust anchor accepted by the relevant relying party and the paths that its policy permits. Cross-signing may provide an alternative route, create a common dependency, or have no material effect on the particular service path.

The consequence also differs between withdrawal and compromise. A surviving independent path can preserve availability after an anchor is removed. It does not necessarily prevent forgery while the relying party continues to accept a compromised path. A statement that the service has two roots therefore requires a scenario-specific interpretation before it becomes a resilience claim.

Payment trust roots and key horizons

Card payments depend on scheme certificate authorities operated by Visa, Mastercard, American Express, Discover, JCB and UnionPay. Each is a single point for its rail. EMV certificate authority key expiry is scheme- and key-length-specific.

Visa has extended its production 1984-bit CA key to 31 December 2029, following EMVCo's annual RSA key-length assessments. EMVCo records that NIST will not support EMV's longest RSA key length, 1984 bits, after 2030.

Issuer keys minted today will remain on cards when that horizon arrives. EMVCo's public position is that quantum computing is not expected to threaten the EMV infrastructure until at least 2040, while the G7 Cyber Expert Group, NIST and the UK NCSC work to horizons between 2030 and 2035. An institution can therefore have different planning horizons for its card and enterprise estates. It needs to record which horizon governs each service.

Custody and generation provenance

Layer 4 resolves the deployed custody arrangement to its relevant firmware family. A shared manufacturer is additional information, not automatic proof that all of its families share the same code or hardware failure domain. Where a cloud service conceals the underlying module, the provider is an unresolved boundary; its name must not be inserted into a manufacturer or generation-design field as though that resolved the technical identity.

Layer 6 then asks where the key material was generated and which designs support that point. The distinction is important for imported keys and issued estates. It also prevents a generic design class, such as a ring-oscillator source, or a shared certification profile from being treated as proof of one common implementation design.

Custody market and provenance disclosure

ABI Research's payment HSM ranking places Thales first, Futurex second and Utimaco third. Across the broader HSM market, analyst estimates put the top five near four fifths of revenue. These are analyst estimates with limited published methodology: the ordering is reliable, while precise shares are indicative. Revenue concentration is context, not a measurement of the institution's cryptographic failure domains.

Payment HSM vendors publish algorithm support and certification status. Among the vendors surveyed for this paper, none publishes the code lineage of its post-quantum implementation. The published Supplier Lineage Disclosure Request exists to test that claim and, where it is wrong, to correct it.

PCI Security Standards Council listings show a post-quantum support flag, and the entry itself directs readers to the vendor for readiness detail. A validation tests behaviour rather than provenance. Two certified modules can descend from the same upstream implementation.

Applied Quantum is dispatching lineage disclosure requests to the principal payment HSM vendors under the published CCF Supplier Lineage Disclosure Request, with a thirty-day response window and a fourteen-day factual-review extract before anything prints. Answers and refusals alike publish here at v1.0 final, under the instrument's three-state record: an attributed reply, an anonymised aggregate for suppliers electing it, or a named non-response for suppliers taking neither.

Negotiated outcomes and function-specific evidence

Layer 5 records the cryptographic outcome relevant to the function. For a key-establishment path, that may include the protocol, version and negotiated group. For an authentication function, the record must identify the authentication method actually used. A message format, a 3-D Secure version number, or a list of supported cipher suites does not by itself resolve the cryptographic outcome.

In payment environments, fallback behaviour must be established from the selected protocol profile, configuration and observations. It is not safe to assume that every card and terminal combination will fall back through the same authentication sequence. The assessment records the operating behaviour and its evidence rather than substituting a general account of what the technology can support.

Messaging and settlement integration

SWIFT carries the great majority of cross-border traffic, and the principal alternative depends on it. The last public estimate, from 2022, put around 80% of CIPS payments on SWIFT messaging (CSIS, citing Yeung and Goh). No later figure has replaced it. This is a dated estimate, not a current measurement of routing.

The BIS Project Leap Phase 2, with the Bank of Italy, the Bank of France, the Deutsche Bundesbank, Nexi-Colt and Swift, replaced RSA signatures with CRYSTALS-Dilithium at NIST security category 3 on TARGET2 liquidity transfers. The report was published on 11 December 2025. Every functional test passed, but a message with a valid post-quantum signature could not complete settlement because no corresponding certificate existed in the T2 static reference data. This was expected in the test configuration. The report states that the system relies on centralised certificate data under PQC (§4.2, p. 21).

The existing configuration could not validate traditional and post-quantum signatures side by side, so hybrid operation would require substantial evolution of the system (§3.3, p. 18). Verification ran about 7.5 times slower: 209.9 milliseconds against 28.1 for RSA (pp. 21–22). The result concerns the tested implementation and environment, not every production architecture. It demonstrates that migration can fail at the service level without an algorithm being broken.

Entropy sources and generation design

Payment HSMs, terminals and the issued card estate depend on a narrow set of hardware random number generator designs. Secure elements typically use ring-oscillator true random number generators conforming to one certification profile, feeding deterministic generators built to one specification. NXP's SE050 documents an AIS-31 class PTG.2 true generator seeding an AIS-20 class DRG.4 deterministic generator. That is a certification profile: evidence towards the design, not a statement of its identity.

Profile conformance raises the likelihood of a shared design without establishing it. Where design identity is not disclosed, CCF records a reach bound rather than asserting a common design as fact. The distinction lets the assessment show the potential commonality and the evidence needed to resolve it.

4. MEASUREMENTS AND INTERPRETATION

A CCF service report combines concentration, largest share, failure-domain reach, criticality and Effective Coverage. These quantities describe different aspects of the same service. Their value is in the combined interpretation, not in reducing them to a single security grade.

CCF defines two principal figures: the Cryptographic Concentration Index, a Herfindahl–Hirschman index per layer per important business service in a form familiar to supervisors, and Effective Coverage, the path-complete migration figure reported with classified counts of the counterparties gating the shortfall. Failure-domain reach explains commonality beyond the index's exclusive assignments.

Three design choices govern their interpretation: asset-count weighting with a flagged-only compensating view under five objective tests; non-disclosure represented in a reach bound rather than as a score change; and layer 1 reported separately as the migration baseline. Layer 1 remains at 10,000 until the last quantum-vulnerable public-key asset retires. That result establishes the limit of procurement as a response: another supplier cannot remove the shared algorithm dependency. Only migration moves the remaining population. Excluding layer 1 from the headline keeps changes at the other five layers visible.

Primary-dependency concentration

The Cryptographic Concentration Index uses a Herfindahl-style calculation: square each dependency's fraction of the applicable asset population, add the squared fractions, and multiply by 10,000. One dependency gives 10,000. Two equal dependencies give 5,000. Four equal dependencies give 2,500.

The index uses one primary assignment per asset within a layer. Alternative paths and overlapping upstream memberships are recorded through reach. Where a function requires multiple components simultaneously, including a hybrid construction, the assessment uses a documented composite primary unit rather than giving one asset two votes. Its constituent implementation relationships remain visible in the reach record.

The asset denominator is a methodological choice with limitations. It is not transaction volume, settlement value, revenue or the fraction of a service that would stop. A single critical signing endpoint may count once while a large population of routine endpoints contributes many entries. CCF requires the asset classes and counting rules to be stated and uses a flagged-only view to examine qualifying critical material separately.

A broad bill of materials also contains records that are not additional denominator entries. Algorithm definitions, libraries, path-state records and service metadata describe the assessed assets; counting every such record as an asset would make the result depend on how much explanatory structure a producer included. Canonical identity and de-duplication are therefore part of the method, not a data-cleaning detail left to each analyst.

Largest share and bands

CCF's bands are structural screening categories. They are not regulatory thresholds or empirically calibrated probabilities of failure. The largest dependency share must be

reported beside the index because a band does not identify the shape of the distribution.

For example, two equal dependencies produce a CCI of 5,000 and a largest share of 50%. That exceeds the 4,900 boundary associated with the lower-bound contribution of a 70% dependency, but it does not mean either dependency has a 70% share. The mathematical implication works in one direction: a very large dependency forces a high index; a high index does not uniquely identify that dependency's size.

The operational question remains separate. A dependency supporting 30% of the asset population might be the only one able to authorise settlement. Another supporting 70% might have a tested substitute with sufficient capacity. The index identifies the distribution to investigate; the failure-impact determination establishes the service consequence.

Reach and overlap

Failure-domain reach is the share of applicable assets exposed to a named domain under a stated failure mode. A shared ancestor may reach several executing families. An asset can lie within more than one upstream domain, so these shares are not added together.

Where membership is incomplete, a named domain has an evidenced lower value and a justified upper bound. The lower value includes established memberships. The upper bound includes additional memberships that cannot be excluded within the specified scenario. Neither the bound nor the width of the interval is a probability. A wide interval indicates that the dependency relationship remains materially unresolved.

The distinction is particularly important for generation design. Suppose three identified generation points account for 63.1%, 25.0% and 11.9% of a population and share one certification profile. The profile does not establish one common design. The bounded design node has known reach at the largest point's own design, 63.1%, and an upper bound of 100%. The lower value is a floor by construction, not evidence that two points share a design.

An evidence request can narrow that bound without changing the deployed estate. It may establish independent designs, confirm a common design or show that the original point-level assumption needs refinement. The report should credit the improvement in knowledge without describing it as a physical diversification that has already occurred.

At layer 2 in the worked example, an ancestor has 57.1% evidenced reach and a 78.6% upper bound. The 21.4 reach points concern 900 assets whose lineage is undisclosed. The executing estate has not changed because the supplier did not answer; the uncertainty is in the reach of its possible ancestor. That spread gives procurement a specific evidence requirement to put to the supplier.

Criticality and asset-count weighting

The five CCF tests flag material that certifies, derives or wraps other keys; authorises an irreversible or finality-bearing action; cannot be revoked within tolerance; remains valid beyond the applicable migration horizon; or lacks an operational equivalent. The service or key owner supplies the facts, and the assessment owner records and challenges the determination.

The flagged-only index is calculated over qualifying assets, not over every asset belonging to a dependency that happens to carry one flagged item. This avoids converting a small critical subset into an inflated population. The flagged count is always reported because small subsets naturally produce high concentration readings.

The same underlying asset may have different criticality in different services. A replacement arrangement may sustain one service but not another, and their tolerances may differ. A single shared CBOM record cannot be assumed to carry every service-specific determination without an explicit contextual record.

5. EVIDENCE ARCHITECTURE

The framework is a suite of documents because inventory production, risk assessment and executive reporting are different activities. Requiring every reader to work through a single large specification would obscure the hand-offs on which a repeatable assessment depends.

The **Universal Framework** defines the method: service scoping, dependency resolution, computation, failure-impact determination and reporting. The **CBOM Profile** defines the data properties and their relationship to CycloneDX. The **Conformance Statement** specifies which inputs are needed for the CCF outputs being claimed. The **Payments Extension** identifies sector-specific populations and issues without changing the core arithmetic. The **Technical Companion** explains how engineering teams can obtain and maintain the evidence. The instruments support requests, determinations, supplier engagement, validation and reporting.

These documents should be used together, but they do not all have the same authority. The Profile controls field names and value semantics. The Universal Framework controls the assessment method. An informative whitepaper explains the rationale; it does not create an additional conformance requirement through an illustrative example.

Profile and base-model relationship

The CycloneDX specification provides native structures for components, cryptographic assets, dependencies and pedigree. The Applied Quantum CBOM Profile uses those structures and adds domain-specific properties for migration governance, custody, service mapping, path observations and concentration evidence.

Implementation ancestry belongs on the implementation component. An abstract algorithm record describing ML-KEM is not the library that implements it and should not acquire a source-code pedigree as though the two were interchangeable. Likewise, a `provides` relationship can identify a capability supplied by a component without proving that a particular endpoint executed that component during the observation period. The assessment needs the deployment context as well as the structural record.

The Profile is not, by itself, a complete serialisation of every CCF determination. Its current flat properties do not encode the entire ordered-hop graph, every asset-to-path relationship, generation origin, all service-specific criticality decisions or a per-claim evidence record. The data contract therefore requires a linked **assessment supplement** where those facts are necessary. The supplement is not a new collection of invented `appliedquantum:*` fields; it is an explicit record of the additional context needed to reproduce the calculation.

Discovery coverage and completeness

A discovery tool can report high coverage of the objects it is capable of observing while missing whole asset classes. A certificate management system may provide a strong certificate inventory without enumerating application-embedded keys or historical generation provenance. The coverage statement must identify its denominator, observable perimeter and unknown classes.

Where a class has a known population but none of its objects have been observed, observed coverage can be zero. Where the population itself is unknown, the assessment cannot turn that absence into a known zero population. A tool-scoped fraction must not be promoted to a whole-service coverage claim without evidence supporting the broader denominator.

This is why conformance and completeness are different. A CCF-1 dataset can contain the fields needed for several indices over a clearly defined subset. That does not make it a complete inventory of the service. A higher conformance level adds data capabilities; it does not certify that discovery is exhaustive or that the service is secure.

Claim-specific evidence

CCF retains five evidence basis labels: runtime observed, binary confirmed, inventory declared, vendor attested and inferred. They are not a universal numerical ranking. Runtime evidence can establish a negotiated exchange without establishing the ancestry of the code that implemented it. A supplier attestation may be the most direct available source for a proprietary firmware-family identity.

The Profile provides an asset-level default and more specific lineage and entropy evidence properties. Where these cannot express differing evidence for individual claims, the assessment supplement records the exceptions. A vendor statement

corroborated by a security policy does not become an inventory declaration simply to assign it a supposedly higher tier. The basis, corroboration and limitations should remain separately visible.

This discipline also applies to missing values. An ungraded supplied value may be provisionally treated as inferred and challenged. A value that was not supplied is still missing. Assigning it an evidence tier does not create a fact on which a computation can rely.

Assessment inputs and conformance

A complete inventory is not required before the assessment can begin. Three conformance levels describe what can be computed. At the first level, standard discovery tooling plus service mapping computes three of the six layers: algorithm, trust root and protocol. An assessment reporting those three layers over its stated population is completed work, not a failed attempt at all six. The layer-2 index can also be computed where the executing-family identity is available; ancestry is a separate evidence requirement.

Data availability and assessment quality are reported separately. A figure over 40% discovery coverage is a different object from one over 95%, regardless of whether both can be calculated. Conformance identifies the available data classes; it does not establish that the assessment was well governed or that the whole estate was discovered.

The instruments needed to begin are available. The Data Request Specification states what the internal inventory response must carry. The Supplier Lineage Disclosure Request supplies a grading scheme that distinguishes provenance disclosure from certification cited in place of provenance. Those records give the assessment a defined route for obtaining and challenging the inputs. The related PQC Migration Framework is published separately.

6. END-TO-END READINESS

At any moment, a payment network contains a distribution of cryptographic states: classical systems, partially migrated systems, hybrid deployments, systems marked complete but unverified, and systems that were never inventoried. Different states at neighbouring institutions do not by themselves provide end-to-end protection. A quantum-safe node continuing to transact with a classical counterparty negotiates down, and that path remains harvestable. The exposure is concentrated in the connections between institutions.

An endpoint migration percentage therefore overstates the institution's position when it is used as a measure of end-to-end protection. Effective Coverage is the share of

cryptographic assets for which every party and every hop on the relevant paths operates at target state. An institution reporting 70% migration can have single-digit Effective Coverage on cross-institutional paths.

A migration programme can accurately report that an endpoint supports a new algorithm while the service remains below target on a path that endpoint uses. Capability, configuration and observed operating state are separate facts. Effective Coverage measures the latter against the service's approved target definition.

For a service requiring post-quantum key establishment, the assessment must establish the institution's own state, the relevant counterparty state and the behaviour of every required hop and fallback. A configured offer of a hybrid group is not sufficient where the observed connection uses a classical group. Equally, a service may reject an incompatible connection rather than downgrade. The operating consequence in that case is a connectivity or availability issue, not evidence that a classical exchange took place.

The target must name the cryptographic function. A protected key-establishment exchange does not establish post-quantum authentication. A migration of certificate issuance does not automatically protect every archived signature or message-signing workflow. Organisations may report separate function-specific views and a combined view where all required functions are satisfied, provided the relationship between those views is explicit.

Asset and path coverage

An assessed asset may support more than one relevant path. If one required path remains below target, the asset does not receive full end-to-end credit merely because another path is protected. This is the reason for CCF's minimum-across-paths rule at the identified-asset level.

The rule depends on an actual path map. It is not an instruction to take the smallest available percentage from several unrelated dashboards. The analyst must know which paths belong to the asset, which functions each path performs, and whether the routes are required, optional, historical or outside the assessment window. Those decisions are recorded so that a reviewer can reproduce the coverage population.

For wholly internal services, the analysis still requires own-state evidence but may have no external counterparties. The counterparty count is then a confirmed zero, not an omitted field. This can identify a service where progress depends primarily on internal engineering and approval rather than external coordination.

Hybrid acceptance and evidence

A hybrid can provide useful protection when the approved target accepts it and the complete construction meets the framework's conditions. Those conditions concern the

accepted post-quantum component and parameters, the security of the composition, and resistance to falling below target on the relevant path. For a dual-signature target, the verifier's actual acceptance rule matters; merely carrying two signatures is not proof that both are required.

Implementation and integration remain relevant. Shared code between hybrid components does not automatically invalidate the cryptographic composition, but it may create a common implementation failure domain. A combiner, parser or verification-policy defect can affect the protection of the complete construction. CCF records those relationships rather than treating the presence of two algorithm names as two independent operational controls.

Where the combiner or fallback behaviour cannot be established, the assessment records `hybrid-unverified` and does not award confirmed target-state credit. That is an assessment status, not an additional value silently inserted into the Profile's migration-status enumeration. An unverified construction is not asserted to be broken; its claimed protection has not been established for the assessment.

Populations and joint-state evidence

Payment services often depend on issued or acceptance-device populations. A population percentage is useful for programme management, but separate percentages do not reveal whether the same members meet each required condition.

Suppose 80% of one defined population satisfies an endpoint requirement and 70% satisfies a second required condition. Without information about their overlap, the joint proportion can lie between 50% and 70%. Reporting the minimum, 70%, as exact coverage would assume that every member satisfying the second condition also satisfies the first. That nesting may be true, but it requires evidence.

The method therefore uses identified joint segments where available and feasible bounds where only comparable marginals are known. Different populations cannot be combined using that calculation without first establishing a common denominator. This distinction prevents a mature endpoint programme and a mature counterparty programme from being reported as a mature end-to-end service when their remaining gaps affect different members.

Counterparty classes and remediation

CCF distinguishes blocking, contractual, peer and population relationships by the practical route to remediation. A blocking party requires coordination or industry engagement; a contractual relationship may permit requirements to be enforced through an agreement or onboarding process; a peer requires bilateral planning; a population typically requires a certification, deployment or replacement programme.

These are contextual classifications. A certificate provider is not automatically blocking in every service, and a large supplier is not automatically beyond contractual influence. The classification should reflect the institution's actual rights, alternatives and operating arrangements. Its purpose is to identify an owner and a feasible next action, not to create a permanent label for a company.

Coverage and coordination in the example

In the processor example in Section 7, Effective Coverage is 0% because no path operates at target state end to end. We expect that result in most payment institutions today; it remains a result for each institution to establish from its own mapped paths rather than infer from another institution's assessment.

The counterparties explain the work behind the shortfall: 6 blocking parties, 12 peers, 340 contractual relationships and a terminal estate at 0% target state. The blocking six are four card schemes and two central bank rails that cannot be compelled. The twelve peers require bilateral outreach. The 340 contractual relationships form a certification programme. The terminal population requires a refresh campaign. These are four problems with different timescales and owners; a single count of 358 would obscure them.

7. WORKED PAYMENT-SERVICE ASSESSMENT

A mid-size acquiring processor assesses card authorisation over 4,200 cryptographic assets, 180 of them flagged. The illustration is synthetic. The service includes the processor-operated settlement-submission leg, which explains the central-bank relationships in the counterparty register. The full unit counts, service boundary, counterparty classification and tolerance reasoning are published in the Payments Extension; the reference implementation computes each figure from those counts.

Layer 1 is the declared algorithm baseline, excluded from the reported service figure. The table distinguishes the largest failure domain in each layer from its largest upstream node. A standalone primary unit remains a failure domain even where no upstream node is named.

Layer	CCI	Largest share	Band	Flagged-only	Largest failure domain (bound)	Largest upstream node (bound)	Tolerance
1 Algorithm *(baseline, excluded from the reported figure)*	10,000	100%	Single-source-equivalent	10,000	–	–	Fail
2 Lineage	4,184	57.1%	Concentrated	–	57.1% evid. · 78.6% bound (reach node)	57.1% evid. · 78.6% bound	Fail, on the bound

Layer	CCI	Largest share	Band	Flagged-only	Largest failure domain (bound)	Largest upstream node (bound)	Tolerance
3 Trust root	2,370	31.0%	Diversified	7,222	31.0% evid. (flagged 83.3%) (standalone unit)	31.0% evid. (flagged 83.3%)	Fail, flagged
4 Custody	5,726	69.0%	Highly concentrated	10,000	69.0% evid. (standalone unit)	–	Fail
5 Protocol	4,660	61.9%	Concentrated	–	61.9% evid. (standalone unit)	–	Pass
6 Key generation	4,748	63.1%	Concentrated	10,000	63.1% evid. · 100.0% bound (reach node)	63.1% evid. · 100.0% bound	Fail, on the bound

Tolerance results come from the failure-impact determination in Universal Framework C.5, applied to reach-defined domains. Layer 2 fails on the upper bound: emergency patch-and-rotate across the 3,300-asset bounded domain exceeds the two-hour tolerance, while the 2,400-asset evidenced domain passes. Layer 3 fails on the flagged view because revoke-and-reissue under the single authority anchoring the issuer master keys exceeds the same clock. Layer 4 fails on the flagged view outright. Layer 6 fails against the consequence threshold on the design bound: re-keying the issued estate is a multi-year campaign, and keys already generated remain exposed.

A dash in the flagged-only column means the illustration places no flagged material at that layer. A dash in the upstream-node column means no upstream node is named. Layer 1's failure domain remains separate from the structural companion across layers 2 to 6. No absent upstream node at layers 4 or 5 removes the standalone primary unit from the largest-failure-domain result.

The **service CCI is 5,726**, the maximum across layers 2 to 6, at custody. Its largest failure-domain reach is instead the layer-6 design bound at 100%. The two figures answer different questions.

Critical subset and custody

Trust root is the only layer that appears diversified overall, at CCI 2,370. Over flagged material, it reads 7,222 because issuer master keys concentrate under one authority. The largest flagged share is 83.3%. An assessment reporting only the headline trust-root figure would miss that concentration.

At custody, the full population gives CCI 5,726, while the flagged-only result is 10,000. All 180 flagged assets are on the first firmware family. This identifies the dependency to examine when testing recovery of critical functions; it does not by itself establish that the firmware is defective.

Ancestry bounds and the tolerance boundary

At layer 2, the ancestor has 57.1% evidenced reach and a 78.6% upper bound. The three executing families remain distinct in the CCI, which stays at 4,184. The uncertain membership concerns a 900-asset family; the other 900-asset family is excluded on the example's evidence.

The tolerance boundary lies inside the spread. The evidenced 2,400-asset domain passes the two-hour clock, while the 3,300-asset bounded domain fails it. The 21.4 reach points make one supplier's silence relevant to a specific operational conclusion. Procurement can request the evidence needed to resolve the membership while engineering tests the recovery constraints. The result is a failure of the upper-bound scenario, not a claim that the full shared membership has already been proved.

Generation-design bound

Three generation points share one certification profile without a disclosed design. A single design behind the whole population cannot be excluded, so the upper bound is 100%. The 63.1% evidenced figure is the largest point's own design, a floor by construction. It is not evidence that two points share a design. Disclosure can narrow the bound without changing the primary generation-point distribution.

The operational consequence concerns the issued material as well as its generator. Replacing a generator does not repair keys it generated earlier. In this illustration, the design-bound scenario fails the consequence threshold because re-keying the issued estate takes years and the earlier exposure remains.

Coverage and counterparties

Effective Coverage is **0%** because no path in the service operates at target state end to end. This is a computation over the mapped population, not an inference from the CCI or the classical baseline. If path data were missing, the result would be uncomputable rather than a fabricated zero.

The register contains 6 blocking counterparties, 12 peers and 340 contractual relationships, together with a terminal population at 0% target state. The blocking six are four card schemes and two central bank rails within the stated service boundary. The schemes and rails need coordinated technical and operational milestones; the peers need bilateral outreach; the contractual relationships need a certification programme; and the terminal estate needs a refresh campaign. A total of 358 named counterparties would not communicate these differences.

Assessment decisions

The service owner can distinguish critical custody and trust dependencies from uncertain ancestry and generation relationships, and distinguish both from the coverage shortfall.

Critical functions require architecture and recovery decisions. Reach bounds require disclosure and scenario testing. The coverage shortfall requires path-specific migration and counterparty coordination.

No single action resolves all three. Buying another HSM product may leave generation provenance unchanged. Obtaining a lineage disclosure may improve confidence without diversifying the estate. Migrating a negotiated exchange may improve coverage while leaving a shared firmware domain intact. A useful remediation plan states which output it intends to change and why that change improves the service's position.

8. MIGRATION DECISIONS AND INTEGRATION

Post-quantum migration changes an operating system of protocols, implementations, keys, certificates, policies and counterparties. An algorithm can be implemented correctly while the surrounding service is not ready to use it. Concentration assessment helps identify shared constraints, but it does not replace interoperability and performance testing.

Project Leap Phase 2, discussed in Section 3, demonstrates the importance of integration and configuration. A valid signature did not supply the certificate data required for settlement. The measured verification difference is specific to that implementation and environment, not a forecast for every production architecture.

For an institution, the corresponding control question is whether the complete path has been tested under the selected cryptographic profile, including certificate and trust data, message handling, counterparties, fallback policy and operational capacity. A successful algorithm benchmark answers only part of that question. A vendor announcement establishes neither the deployed product's readiness nor the behaviour of the institution's configured path.

Deadlines and applicability

The migration schedule should distinguish binding requirements, supervisory guidance, industry milestones and internal targets. These sources can influence one another without becoming legally equivalent.

Discussion centred on Q-Day invites delay because the date is uncertain. The binding migration clock is nearer and fixed. The instruments contributing to that clock must nevertheless be read by their actual scope and status.

NIST's transition roadmap, IR 8547 – still an initial public draft at this edition – would deprecate quantum-vulnerable public-key algorithms at 112-bit security strength, RSA-2048 among them, after 2030 and disallow every quantum-vulnerable public-key algorithm after 2035. US national security systems must be quantum-safe by 2035.

The UK NCSC roadmap sets discovery by 2028, priority migration by 2031 and completion by 2035. These are migration milestones, not predictions of when a quantum attack becomes practical.

The United States issued Executive Order 14412 on 22 June 2026, with publication in Federal Register Vol. 91 No. 121 on 25 June 2026. It requires every federal high-value asset and high-impact system to reach post-quantum key establishment by 31 December 2030 and post-quantum digital signatures by 31 December 2031. Section 5(d) directs FAR rulemaking to bind covered contractors to the same FIPS requirements by end-2030. That is directed rulemaking, not an accomplished contractor mandate; our analysis explains the distinction.

The primitive-specific deadlines separate decryption and forgery in the federal programme. Procurement then transmits the timetable: a federal deadline becomes a supplier deadline and affects the other institutions buying the same technology.

The G7 Cyber Expert Group statement of January 2026, co-chaired by the US Treasury and the Bank of England, sets a coordinated direction for the financial sector. Europol's Quantum Safe Financial Forum has pressed the case since February 2025. Few of these instruments bind a commercial payment institution directly. Institutions nevertheless inherit the horizon through certified products, shared vendors and supervisors' expectations.

Payment schemes and technical standards can introduce additional product and estate constraints. The EMVCo discussion of elliptic-curve support concerns a distinct evolution of the EMV cryptographic profile. Moving from RSA to elliptic-curve cryptography does not remove the Shor-type public-key dependency addressed by CCF layer 1. An institution needs to record which requirement, function and estate a milestone concerns rather than combine all cryptographic changes into one migration date.

The CBOM Profile's horizon fields support that record, but the field value is not the legal analysis. A service owner needs a source, status, jurisdiction, affected function and applicability determination. An internal target may be more conservative than an external requirement; it cannot make a binding deadline disappear by choosing a later date.

Migration compression and implementation assurance

The deadline generates correlated implementation risk as well as reducing exposure to the algorithms being retired. The sector migrates towards one date using a small set of libraries under the same time pressure. That is the condition under which rushed implementation failures cluster rather than remain isolated.

In June 2026, Daniel J. Bernstein published an analysis demonstrating recovery of an equivalent ML-DSA signing key in under a second on a single laptop core by exploiting implementation defects. The analysis traces one bug class across three production

implementations. Its estimate that roughly a quarter of ML-DSA implementations ship with a severe vulnerability at first release is a model projection from the historical CVE record under stated assumptions, not an observed census.

Two programme consequences follow. Deferral compresses and synchronises the eventual migration and raises the chance that an institution deploys a library's first version. And "we will adopt the standards" is not a sufficient programme statement. Verification and testing belong beside migration: a wrongly implemented post-quantum node is trusted and cannot be quarantined in the same way as a visibly non-compliant node.

Diversification and operating costs

Maintaining several cryptographic implementations can introduce additional configuration, validation, patching and incident-response work. Independence is not the only selection criterion. An immature or poorly maintained alternative may create a worse overall security position even if it lowers the concentration index.

CCF therefore does not prescribe diversity at any cost. It asks the institution to identify shared domains, understand their consequences and evaluate feasible alternatives. In some services, a well-assured implementation with a tested replacement procedure may be preferable to nominal diversity that the operating team cannot maintain reliably. Where the market offers no credible independent alternative, evidence, recovery design, contractual change notification and explicit risk acceptance may be the appropriate controls.

The same reasoning applies during migration. A temporary increase in implementation concentration may accompany a justified move to a better-supported target platform. The decision should state the trade-off, the service consequences and the planned response. A metric should make that decision more transparent, not force engineers to optimise a number independently of security and operability.

9. ACCOUNTABILITY AND CONTROL EFFECTIVENESS

A high CCI is an observation about dependency distribution. A decision requires a failure scenario, affected functions, service effect, recovery assessment and consequence threshold. This is the role of the failure-impact determination, which is the link between the measurement and the institution's operational resilience framework.

The service owner supplies operational facts. Those include the time to detect the problem, decide on the response, obtain replacement capability, deploy or redistribute it, and restore the required service. A nominal recovery-time objective or an untested supplier statement is not automatically a demonstrated recoverability estimate. The assessment owner records the evidence and challenges the assumptions.

Confidentiality and integrity require additional care. A service may remain available while protected data is exposed or forged instructions can be accepted. An availability tolerance alone does not determine whether those consequences are acceptable. The institution must define the relevant consequence threshold and test it explicitly. Missing thresholds are governance findings, not implicit permission to record a Pass.

The resulting outcomes are Pass, Fail or Not assessable. A Fail based on an upper-bound membership scenario remains labelled as such. A Not assessable result identifies missing operational facts or thresholds; it is not a low-risk category. These distinctions allow a risk committee to tell a demonstrated tolerance breach from a scenario-sensitive conclusion and from a gap that prevents a conclusion.

Control effectiveness

Payment institutions operate reconciliation, dual controls, limits, manual fallback and kill switches. The assessment concerns residual rather than wholly uncontrolled exposure. These controls have different effects across the harm channels.

Control	What it contains	Where it fails here
Reconciliation	Forged or erroneous entries after settlement, cutting fraud loss	Acts after the fact, trusts its own reference data, and is swamped at common-mode scale. Does nothing for data already harvested.
Dual controls	Single-actor authorisation abuse	Both approvers rely on the same cryptographic trust. Four eyes verify the same forged signature.
Limits and velocity caps	Per-transaction and per-counterparty loss	A common-mode event breaches many limits at once. Limits neither restart halted rails nor recover exfiltrated data.
Manual fallback	Some throughput during systems failure	Runs at a fraction of volume, which is how an outage becomes a multi-day backlog. Manual identity checks are themselves compromised in a trust failure.

Control	What it contains	Where it fails here
Kill switches	Contagion from a compromised connection	Invoking them is the availability loss. They require identifying the compromised node, which an invisible implementation flaw denies.
Netting and central clearing	The direct-default cascade	Does not touch the trust-and-availability channel, where the systemic weight sits.

Containment is strongest against authentication and fraud, the channel hardened for decades and the smaller part of this exposure. It is weakest against availability, where operating the controls that help is itself part of the disruption, and it does almost nothing for confidentiality once data has been harvested. Residual exposure is lower than inherent exposure, but the reduction is concentrated in the channel that matters least in this analysis.

An assessment still records each control's expected contribution. An independent reconciliation source may help identify invalid instructions, while one relying on the same compromised trust path may not provide the intended independence. A manual fallback may preserve a critical subset of payments without sustaining normal throughput. Isolating a connection may contain exposure while creating an availability shortfall. Record the protected property, operating conditions, capacity and evidence before assigning the control a benefit.

Long-tail dependencies

The likeliest point of failure is the long tail of small, under-resourced providers and institutions that migrate slowest. It is also the population most often missed by an incomplete inventory. Migration protects an institution from its own compromise; it does not preserve trust in counterparties that failed to migrate.

At an algorithm break, institutions still using the affected algorithm fail at the same time. A migrated counterparty cannot distinguish valid from forged messages across that population, so its rational response is to suspend the population's connections. The systemic consequence is interrupted settlement, rather than bankruptcy. Across a fragmented tail, third-party assurance thins to a questionnaire. The probability of failure is highest where control is lowest.

Within an institution, the exposure spans cryptography, third-party risk, operational resilience, technology and treasury. It is therefore often left without a single accountable owner, despite each function owning part of the underlying work.

Responsibilities and decision rights

CCF places assessment ownership in the second line. First-line engineering and service teams produce inventory and operational evidence; second line challenges completeness, resolves assessment treatment and presents the results to the accountable decision-maker. Independent validation examines the computation and material judgements according to the institution's governance arrangements. Internal audit may provide independent assurance, but it should not be treated as the default owner of model development or routine assessment validation.

The instruments support this division of responsibility. The Data Request states the required population and evidence. The Supplier Lineage Request obtains information that internal discovery cannot establish. The Determination Log records judgements and their challenge history. The Model Documentation Pack records classification, assumptions, limitations and validation. The Board Report presents material decisions without detaching the institution-level figure from its qualifiers.

A supplier response must also be recorded accurately. An explicit refusal, silence after the deadline and a received answer that does not address the question are different facts. The disclosure response rate counts full and partial responses under the defined request scope, but it is not the proportion of the asset estate with established ancestry. Both the response record and the remaining technical gaps matter.

10. REPORTING AND GOVERNANCE

The primary CCF output is a service-by-layer heat map supported by the determination record. It shows the concentration index, largest share, critical subset, relevant reach and tolerance outcome. Effective Coverage and the counterparty register appear alongside the service, with the evidence and population qualifications needed to interpret them.

An institution-level indicator can be useful for oversight, but it compresses information. CCF defines it as a service-exposure-weighted mean of the service maxima across layers 2 to 6. The figure is withheld below assessment maturity M3, where the required governance and challenge conditions have not been established. Meeting the maturity gate does not excuse missing weights, unresolved scope or incomplete reporting qualifications.

Where some layers have not been assessed, the service figure is labelled a **computed-layer maximum**. It is not presented as the maximum of the unmeasured estate. Where a

reported result is bounded, retain the bound and the population basis that supports it. An unknown population is not repaired by assigning every missing item to a convenient extreme.

The institution-level figure is accompanied by the count of services at or above 8,100, the lowest Effective Coverage with any necessary qualification, the `NOT MAPPED` service count, the supplier disclosure response rate where requests were issued, and the largest structural failure-domain reach across layers 2 to 6. Layer 1 remains a separate baseline so it does not saturate the structural companion every time classical public-key material is present. Not-assessed layers and Not assessable tolerance results are also made visible.

Governance mappings

Cryptographic concentration crosses existing risk and control regimes. The following mapping identifies the relevant integration points and the additional dependency evidence needed.

Regime	Where it maps, and what is missing
Risk taxonomy, RCSA and appetite	Hidden under generic technology or cyber risk. Needs its own line as a correlated concentration exposure, with an appetite statement naming shared-cryptography concentration.
Operational resilience (PRA SS1/21, Basel principles)	Map cryptographic failure to important business services and test a common-mode scenario against impact tolerances. Few severe-but-plausible scenario libraries yet include one.
Third-party and concentration (DORA Art. 28–30, UK CTP)	The closest fit, and its substitutability test misses shared-lineage concentration. Add algorithm, lineage, PKI and firmware overlap to the concentration assessment and the register.
FMI and cyber resilience (CPMI-IOSCO PFMI Principle 17)	Settlement finality depends on cryptographic integrity. Record that dependency and fold a crypto-trust scenario into resilience testing.
Cryptography and PQC (FIPS 203/204/205, IR 8547, PCI DSS 12.3.3)	The programme must add complete cryptographic inventory, crypto-agility and implementation verification, not only

Regime	Where it maps, and what is missing
	algorithm adoption.
Systemic importance (Basel G-SIB)	The interconnectedness and substitutability indicators already capture the principle. Scored on that ruler, shared cryptographic providers would rank as systemic.
Model risk (SR 11-7, PRA SS1/23)	Whether computing this metric constitutes a model is the institution's documented determination. Either way it includes documented assumptions, stated limitations and independent review.

Systemic-importance proposition

The Basel G-SIB assessment scores institutions on interconnectedness, substitutability and footprint to identify nodes too central to fail. Our proposition is that applying the same indicators to the shared layers beneath payments would rank the largest node in each layer above the most systemically important individual bank. Those nodes are not banks, hold no capital against that weight and are not scored on that basis.

This is a proposition awaiting computation on real exposure data, not a completed G-SIB comparison. The dependency evidence collected through CCF makes it testable.

Measurement changes

Improved discovery can raise or lower a concentration reading. If it finds more assets on the largest existing dependency, the index may rise. If it identifies previously unseen independent units, the index may fall. Neither direction alone establishes an improvement or deterioration in the deployed estate.

CCF therefore separates estate changes from measurement changes. Deployment, retirement and re-platforming are distinguished from discovery, scope change and correction of attribution. The counterfactual calculation in Universal C.8 provides a fixed accounting convention for reporting their contributions. It is a reproducible decomposition, not a claim that nonlinear changes have one uniquely determined causal allocation.

The same distinction applies to reach. A new supplier disclosure may narrow an interval without changing any asset. Retiring a population may reduce an upper bound without providing any new ancestry evidence. Reporting both as undifferentiated risk reduction

would create the wrong incentives for the inventory, procurement and engineering teams.

A board report should therefore identify what changed, why it changed and which action is responsible. It should also state where a prior population cannot be reconstructed reliably enough to support the comparison. A missing historical denominator is a comparability limitation, not a reason to label the unexplained residual as remediation.

11. FIRST ASSESSMENT

The first cycle should be designed around a decision that the institution can act on. A migration sequencing decision, a custody renewal, a service resilience review or a supplier evidence programme gives the assessment a practical purpose. A mandate to produce an enterprise score without a defined decision tends to put pressure on scope and data quality before either is established.

A ninety-day cycle produces a first result. Name the assessment owner in second line and select five to eight services. Select for materiality and coverage of the cryptographic estate, not solely because inventories are convenient. A narrow, well-qualified assessment is more useful than a broad one that silently excludes difficult dependencies.

The first phase establishes the owner, boundary, asset classes, tolerances and target-state definitions. Request existing inventory data and issue supplier disclosure requests in week one. Resolve the computable layers and report with the discovery-coverage fraction attached. The full first-cycle specification is in the Universal Framework, Part B.

CCF-1 supports the indices for layers 1, 2, 3 and 5 where their required identities and operating evidence are available. CCF-2 adds custody and criticality inputs. CCF-3 adds generation-point and advanced lineage/design evidence for the fuller assessment. These are data-capability claims over the stated scope, not security grades. A complete path assessment has its own EC marker, independent of the CCF level.

The Technical Companion provides the engineering route: reconcile inventory sources, establish executing components, construct accepted certificate paths, inspect custody and generation records, obtain supplier evidence, and capture negotiated outcomes. Ninety days will not produce lineage and entropy evidence at usable tiers. Report the disclosure response rate instead: it is the fair measure of an institution whose constraint is supplier refusal rather than failure to request the evidence.

The first report should include actionable findings and explicit gaps. A confirmed common domain may justify remediation immediately. A materially wide bound may justify targeted disclosure or validation. A missing consequence threshold should be assigned to the service and resilience owners. An assessment should not be held

indefinitely for a higher conformance label when its current findings can already inform decisions.

12. SECTOR USE AND AGGREGATION LIMITS

A shared dependency can affect several institutions, making sector-level analysis useful. However, adding institution-level CCI values does not reveal a sector failure domain. The analysis requires common dependency identities, compatible service definitions and a method for resolving overlap across submissions.

A supervisor or industry body would need to distinguish a measured deployment population from a supplier-market survey. A vendor revenue share is not an implementation-lineage share. A count of implementations in open-source repositories is not a count of independent families deployed in financial services. The count of independent post-quantum lineages is the census proposition in Section 2, not an established market-wide count. The example does not establish the share of payment assets generated by one design.

Sector contributions also contain sensitive operational information. An aggregation arrangement should specify the minimum fields needed, permitted uses, access controls, treatment of supplier-confidential evidence and publication rules. Public reporting should not imply that a named supplier is vulnerable merely because it appears in a common dependency domain or declined a disclosure request.

The separate Sensitivity Model Concept Note explores impairment propagation through a specified network. Its topology is assumed. It also needs weights, service transmission rules and response assumptions in addition to measured CCF inputs. Its proposed outputs describe ordering and thresholds rather than measurement, and no monetary figure should be derived from them. Any threshold is conditional on the stated intervention and model configuration, not an inherent property of a CCI value.

The immediate value of CCF does not depend on that additional modelling. A service owner can act on an evidenced common implementation, a critical custody dependency, an unresolved generation origin or a below-target counterparty path without first estimating system-wide contagion.

Measured and assumed content

Every concentration and lineage claim in the incident and payments-stack sections cites a public primary source inline. Market-share figures are analyst estimates with limited published methodology; their ordering is more reliable than their precise percentages. The CIPS routing figure is a 2022 estimate and is expressly dated. The G-SIB comparison is a proposition pending computation on real exposure data. The worked example is illustrative, not an assessment of an actual institution.

The sensitivity concept note is published separately because its assumed network and prospective outputs are distinct from these structural measurements. The argument in this paper does not depend on that model.

Validation against the public record

The Universal Framework's Annex 1 re-runs the layer-6 reading over the 2017 ROCA disclosure: one Infineon generation design behind a brand-diverse estate, certification records as the pre-disclosure evidence, and a re-keying response expressed in Estonia's suspension of 760,000 ID-card certificates in one government decision. The structural reading was computable from evidence available before disclosure. The claim is about identifying the common dependency and its response implications, not predicting the cryptanalytic flaw.

13. ASSESSMENT DECISIONS

Each material finding should produce a decision with an owner, evidence requirement and review date. Six actions establish that discipline.

1. **Long-tail ownership.** Assign one accountable executive across cryptographic inventory, dependency concentration and long-tail migration. Execution remains across the three lines; accountable ownership must not disappear between their responsibilities.
2. **Substitutability assessment.** Add algorithm, lineage, PKI and firmware overlap to third-party concentration assessment and the register under DORA Article 29. Supplier identity alone does not establish an independent cryptographic failure domain.
3. **Lineage disclosure.** Make provenance a procurement and due-diligence requirement, with contractual notification of change. Treat certification as necessary and insufficient evidence for this purpose.
4. **Inventory and visibility.** Build the cryptographic inventory and record the classes and populations it cannot see. Preserve those limitations wherever the assessment is reported.
5. **Implementation verification.** Test implementations rather than only adopting algorithm standards. Prefer sound hybrid constructions as defence in depth: a hybrid containing a classical component reduces the consequence of a break, not the concentration. Assure the deepest dependency layers first.
6. **Measurement and governance.** Compute the index and Effective Coverage on the institution's own estate. Set an appetite threshold from its own distribution, make and record the model-classification determination, and govern the computation accordingly. Complete migration ahead of the mandate rather than at it.

These actions do not require the institution to eliminate every common component. Some concentration follows from interoperability, standardisation and the available supply of supported technology. The governance obligation is to understand the dependency, establish the consequence and maintain a response proportionate to the service.

CONCLUSION

Shared cryptography makes the payment network a correlated dependency position. The post-quantum transition will test that position, and the deadline intended to manage the transition is itself a source of correlated implementation risk.

The algorithm standards exist. CCF supplies the operational assessment of inventory, provenance and coverage needed to connect those standards to service decisions. It distinguishes the dependency that requires redesign, the missing evidence that requires disclosure and the path that requires counterparty coordination.

The constraints are time, coordination and ownership. An institution needs to know which shared failure it can withstand, which it cannot, and who is responsible for closing each remaining gap.

SELECTED SOURCES

Primary sources are linked inline throughout. Additional references:

- Basel Committee on Banking Supervision. *Global systemically important banks: assessment methodology* (2013, rev. 2018).
- Battiston et al. *DebtRank.* Scientific Reports 2, 541 (2012).
- Eisenberg & Noe. *Systemic Risk in Financial Systems.* Management Science 47(2), 2001.
- Eisenbach, Kovner & Lee. *Cyber Risk and the U.S. Financial System: A Pre-Mortem Analysis.* JFE 145(3), 2022.
- Glasserman & Young. *Contagion in Financial Networks.* J. Economic Literature 54(3), 2016.
- Haldane. *Rethinking the Financial Network.* Bank of England, 2009.
- CPMI-IOSCO. *Principles for Financial Market Infrastructures*, 2012.
- Mosca. *Cybersecurity in an Era with Quantum Computers.* IEEE S&P 16(5), 2018.
- NIST. FIPS 203/204/205 (2024); IR 8547 initial public draft.
- ASC X9. *Post-Quantum Cryptography Financial Readiness Needs Assessment*, 2025.
- FS-ISAC. *The Impact of Quantum Computing on the Payment Card Industry*, 2025.
- CycloneDX / ECMA-424, cryptographic bill of materials specification.

RESOLVED AT V1.0-RC

Three questions carried from review, closed August 2026.

1. Section order stands. The evidence in Section 3 precedes the explanation of the measurement gap in Section 4, concrete before diagnosis.
2. The paper carries the worked example's summary table and selected readings. The full worked example, unit tables and companion set publish in the Payments Extension.
3. The ninety-day start stays, held to its current half page, with the full first-cycle specification in the Universal Framework.