AUGUST 2026

APPLIED QUANTUM

THE APPLIED QUANTUM CRYPTOGRAPHIC CONCENTRATION FRAMEWORK

UNIVERSAL FRAMEWORK

Measuring cryptographic concentration beneath the vendor layer

Version 1.0-RC – August 2026

Steve Vaile and Marin Ivezic, Applied Quantum

Disclosure: Applied Quantum provides post-quantum migration advisory services. This work is published openly under CC BY 4.0 and is computable without engaging 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	Second-line methodology and practitioner framework
Parent document	None – this is the parent document of the CCF family
Intended audience	Risk analysts and second-line functions – operational and technology risk, model risk, vendor and third-party risk – and the compliance officers and programme managers supporting them
Assumed knowledge	Working knowledge of third-party risk management and the institution’s important business services; no cryptographic engineering background is assumed
Scope	The complete method: units and metrics, evidence standards, the eight-phase assessment, determinations, reporting and interpretation.

	Sector enumerations live in the extensions; discovery tooling sits outside the framework, in the Technical Companion.
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	7 September 2026	Pre-promotion consistency pass: retired double-reading wording aligned to the C.3 reach redesign; cloud-custody guidance aligned to the Universal's unresolved-boundary determination. No figure changed unless the changelog says so.
1.0-RC	8 September 2026	Revision of 8 September 2026: prose revised throughout; A.1 dates the first UK CTP designations; C.7 companion 1 "at or above 8,100"; hybrid counting as TV-10 computes it; profile-only design reach stated as a floor; largest failure domain never smaller than the largest unit; one-unit convention for unresolved boundaries; disclosure states aligned to the Profile enum; convergent-ancestry determination corrected; corroboration recorded separately from evidence basis; certificate identity by fingerprint. No figure changed.

HOW TO USE THIS DOCUMENT

This is the parent document of the family. Everything else – the sector extensions, the payments whitepaper, the Conformance Statement, the Technical Companion and the instruments – extends or operationalises what is defined here, and each cites the phase, part or determination it builds on. Read Part A for the method's shape, then work forward. The extensions are enumeration aids, not replacements for the phases.

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

Purpose and use 8

Part A – Foundations 9

A.1 What CCF measures 9

A.2 Intended use and document authority 10

A.3 Scope and limitations 11

A.4 The six layers 12

Layer 1 – Algorithm 13

Layer 2 – Implementation lineage 14

Layer 3 – Trust root 15

Layer 4 – Key custody platform 15

Layer 5 – Protocol and negotiation 16

Layer 6 – Key generation 17

A.5 Roles and independence 17

A.6 Evidence standards 19

A.7 Determinations and uncertainty 22

Part B – Assessment method 24

Phase 0 – Mandate and scope 24

Phase 1 – Service selection and boundary 25

Phase 2 – Data request and receipt 25

Phase 3 – Resolution 27

Phase 4 – Computation 29

Phase 5 – Coverage and counterparties 29

Phase 6 – Interpretation and challenge 30

Phase 7 – Reporting, remediation and re-assessment 30

Part C – Metric specification 32

C.1 Population, share and weighting 32

C.1.1 Layer applicability and denominators 33

C.1.2 Exclusive primary assignment 34

C.1.3 Unresolved primary identities 34

C.2 Criticality and the flagged-only view 34

C.3 Index and failure-domain reach 35

C.3.1 Cryptographic Concentration Index 35
C.3.2 Failure-domain reach and its bounds 36
C.4 Bands and largest-share interpretation 37
C.5 Failure-impact determination and tolerance test 38
C.5.1 Required record and selection of domains 38
C.5.2 Layer-specific scenario semantics 39
C.5.3 Confidentiality and integrity thresholds 40
C.5.4 Tolerance inheritance and unmapped services 40
C.6 Effective Coverage 41
C.6.1 Calculation and evidence 41
C.6.2 Hybrid acceptance 42
C.6.3 Populations and joint state 42
C.6.4 Counterparty classification and markers 43
C.7 Roll-up and reporting 44
C.7.1 Primary output 44
C.7.2 Service and institution figures 44
C.7.3 Mandatory companions and limitations 45
C.8 Movement between cycles 46
Part D – Interpretation, challenge and materiality 49
D.1 Reading the layers together 49
D.2 Materiality 50
D.3 Challenge questions 51
D.4 Reviewer error catalogue 52
D.5 Feasible diversification 53
Part E – Remediation 54
E.1 Selecting an effective response 54
E.2 Changes that may not address the finding 55
E.3 Sequencing and acceptance 56
Part F – Assessment maturity 57
Part G – Programme integration 59
G.1 Risk taxonomy and RCSA 59
G.2 Risk appetite 59

G.3 Scenario analysis and capital processes 60
G.4 Third-party risk and procurement 60
G.5 Operational resilience 61
G.6 Model risk and independent review 61
Part H – Suggested first-cycle plan 62
Part I – Instruments 63
Part J – Governance, licence and provenance 65
Annex 1 – Validation against the public record: ROCA (2017) 67
The event 67
The retrodiction 68
What this validates 68
Sources, tiered 68
Annex 2 – Canonical asset identity and de-duplication 70
A2.1 Counting principle 70
A2.2 Merging sources 71
A2.3 Identity across assessments 71
A2.4 Division of responsibility with the Profile 72
Resolved at v1.0-RC 72

Sector extensions, aligned on this version baseline: Payments · Financial Services. Published at ccframework.org. Extensions follow the sector taxonomy of the Applied Quantum PQC Migration Framework.

Input data contract. CCF consumes a cryptographic bill of materials conforming to the Applied Quantum CBOM Profile v1.0-RC or later. The CCF CBOM Conformance Statement specifies which fields are required at which level.

Metrics defined here. The **Cryptographic Concentration Index (CCI)**, computed per layer per service. **Effective Coverage (EC)**, its companion.

PURPOSE AND USE

This document defines how to perform a Cryptographic Concentration Framework assessment and interpret its results. It converts cryptographic inventory evidence into a service-level view of primary dependency concentration, shared failure-domain reach and end-to-end target-state coverage. Its intended outcome is an accountable risk decision, not an algorithm catalogue or an unsupported security score.

Read Part A to understand the method and its limits, Part B to organise an assessment, and Part C to implement or review the calculations. Parts D to G cover interpretation, remediation and governance. The sector extensions identify what to enumerate; the Conformance Statement and Profile specify the inputs; the instruments preserve the evidence and decisions through the process.

PART A – FOUNDATIONS

A.1 WHAT CCF MEASURES

DORA Article 29 and the UK Critical Third Parties regime ask whether a provider is substitutable. An institution running four hardware security module vendors, three TLS terminators and two certificate providers can pass that test. Resolving each dependency to the cryptographic components executing underneath can reduce the count to one or two, because every vendor implements RSA and elliptic-curve cryptography, several post-quantum implementations descend from a common ancestor, and some share a hardware random number generator design. Switching vendor changes the supplier relationship and leaves the failure mode intact.

The Cryptographic Concentration Framework (CCF) assesses the extent to which a business service relies on common cryptographic dependencies. Its purpose is to identify failure domains that may remain shared even where suppliers, products and contracts are different. A service using two hardware security module suppliers, for example, may still depend on one firmware family or one key-generation implementation. Conversely, two products from the same supplier may have distinct implementations. Neither conclusion follows from the supplier register alone.

CCF produces three complementary views. The **Cryptographic Concentration Index (CCI)** describes the distribution of assets across mutually exclusive primary dependency units at a particular layer. **Failure-domain reach** describes the assets potentially affected by a specified shared dependency, including relationships that overlap. **Effective Coverage (EC)** describes progress against a service's cryptographic target state across its relevant operating paths, rather than merely at its own endpoints. Criticality and failure-impact determinations connect these structural and coverage measures to service consequences.

The method is defined over cryptographic failure modes. The quantum failure mode is its first instantiated use case because it is the one with a regulatory clock. ROCA, Debian and KyberSlash are classical events at layers 6 and 2 that the layers measure with no modification. Sector and failure-mode extensions beyond post-quantum migration are expected and the Universal Framework changes for none of them. The quantum cryptanalytic threat is the first specified algorithm-layer scenario; implementation, custody, trust and key-generation failures also apply to classical systems. Applying a different algorithm-layer threat model requires an explicit

extension or separately labelled analysis. The present layer-1 grouping must not be reused silently for a different failure mechanism.

CCF supplements existing risk management. DORA Article 29 addresses ICT third-party concentration and substitutability. The UK Critical Third Parties regime took effect on 1 January 2025 and applies to 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. It addresses systemic risks associated with designated providers. These are not limited to considerations that CCF excludes, nor do they prohibit deeper dependency analysis. CCF supplies a specific method for examining shared cryptographic dependencies within those broader assessments; it is not a statement that the regulations require this particular method.

A.2 INTENDED USE AND DOCUMENT AUTHORITY

CCF is a second-line instrument written to be executed by risk analysts, not engineers: the analyst requests data, grades it, challenges it, forms determinations on ambiguous cases and defends the result in review, so the framework specifies data requests, evidence standards and determinations. Discovery procedures, tooling and detection logic are out of scope and sit in the Technical Companion, because a risk framework that specified discovery tooling would age badly and be executed by the wrong function.

CCF is designed for second-line oversight supported by first-line engineering evidence. An assessment should inform an identified decision: risk acceptance, migration sequencing, procurement, resilience testing or remediation funding. A cryptographic inventory without service context cannot by itself answer those questions. Equally, a risk function cannot establish executing code, negotiated outcomes or generation provenance without technical support.

This Universal Framework governs the method, counting rules, calculations and required reporting. The CBOM Profile governs the meaning and carriage of its properties. The CBOM Conformance Statement maps those properties, native CycloneDX structures and the assessment supplement to particular outputs. Sector extensions identify the systems, assets and counterparties to examine. Instruments provide the records through which the method is performed. The Technical Companion, Payments whitepaper and sensitivity concept note are informative.

A conflict is resolved within the relevant domain of authority: this document governs calculations; the Profile governs property semantics. A consumer requirement does not authorise a new property or a changed enumeration in the Profile. Where the Profile does not carry a necessary assessment fact, the fact is retained in the assessment

supplement defined in the Conformance Statement, with a stable reference to the source inventory. It must not be represented by an invented `appliedquantum:*` property.

In normative provisions, **must** denotes a requirement, **should** a recommendation from which a documented departure is permitted, and **may** an option. Examples, suggested schedules and sample risk-appetite statements are informative.

Entry points by reader: the chief risk officer and board read the whitepaper, then C.7, D, F and G. The assessment analyst reads Parts B and C and the instruments. First line and tooling read the Technical Companion, the Profile and the Conformance Statement.

A.3 SCOPE AND LIMITATIONS

The unit of assessment is an important business service, or an explicitly mapped equivalent. The financial-services extensions use that term as an operational assessment unit. Institutions must map it to their own taxonomy: for example, the UK important-business-service concept and DORA's critical or important functions are not interchangeable legal definitions.

The scope includes cryptographic dependency concentration, uncertainty about shared dependencies, target-state coverage, criticality, and the service effects of specified failures. It excludes prediction of a cryptographically relevant quantum computer, estimation of financial loss, certification of cryptographic strength, and management of the migration programme as a whole. The PQC Migration Framework addresses the broader programme; CCF supplies concentration and coverage evidence to it.

The framework excludes the following:

Excluded	Why
Discovery methods and tooling	First-line engineering. See the Technical Companion.
Loss estimation	CCF measures the structure of a dependency, not the money at risk. No monetary figure should be derived from any output.
Timing of a cryptographically relevant quantum computer	Every CCF output is independent of any view on when.
Implementation quality	A dependency with a good security record scores identically to one with a poor record. Concentration and quality are separate questions.

Excluded	Why
Single-product defects	The subject is failure that crosses vendors.
Migration programme management	Covered by the PQC Migration Framework.

A CCI is not a probability of compromise, a measure of implementation quality, a percentage of business volume lost, or a capital charge. High concentration may be unavoidable or appropriately controlled; low concentration may coexist with an intolerable single-point failure. The framework measures a structure and requires a separate determination of its consequences.

CCF focuses on common-mode dependencies, including a single component used repeatedly within one institution. It does not require multiple vendors to be present before a dependency is relevant. A defect affecting only one isolated product may be outside a concentration study, but a single product supporting several critical functions remains within scope.

CCF is written for regulated financial services, where an important-business-service taxonomy and impact tolerances exist. Other sectors need a substitute unit of assessment and a restatement of C.5's tolerance rules. Such use should identify those substitutions rather than imply that the supplied sector enumerations have been validated for every industry.

A.4 THE SIX LAYERS

The primary unit used for each CCI is distinct from the wider failure domains assessed through reach. This separation is fundamental: changing an ancestry record does not necessarily change the executing estate, while a common ancestor can affect several executing families simultaneously.

Layer	Failure mechanism	Primary unit for the CCI	Additional domains assessed through reach
1 – Algorithm	Quantum cryptanalysis of factorisation and discrete-logarithm assumptions	One quantum-vulnerable public-key dependency class	The class itself; protection by a sound hybrid is considered in the consequence determination
2 – Implementation	A defect in executing or inherited	Executing implementation	Relevant ancestors, shared libraries,

Layer	Failure mechanism	Primary unit for the CCI	Additional domains assessed through reach
lineage	cryptographic code	family, or an explicitly defined composite execution unit	components and implementation paths
3 – Trust root	Anchor compromise, distrust, deprecation or a common policy action	Terminating anchor of the primary accepted certification path	Alternative anchors, bridges and policy authorities, under a stated failure mode
4 – Key custody platform	A shared firmware, module or hardware failure	Custody firmware family	Manufacturer or shared hardware/SDK domains where relevant commonality is evidenced or bounded
5 – Protocol and negotiation	Failure of the selected cryptographic profile, including downgrade	Observed outcome tuple for the in-scope cryptographic function	A shared downgrade mechanism or other evidenced common profile dependency
6 – Key generation	A defect in entropy acquisition or key-generation design	Identified generation point	Entropy-source designs, generation algorithms and shared generator components

Layer 1 – Algorithm

Include persistent assets whose in-scope function uses public-key material based on integer factorisation or discrete logarithms. RSA, finite-field Diffie–Hellman and elliptic-curve systems are one dependency class because a fault-tolerant quantum computer running Shor-type algorithms defeats the assumption under all of them. Under the specified quantum scenario, RSA, finite-field Diffie–Hellman and elliptic-curve systems form one dependency class. Distinct parameter sizes are retained in the inventory but do not create separate layer-1 units.

Adequately provisioned symmetric cryptography and post-quantum algorithms are excluded from this class. This exclusion is not a general assurance statement about symmetric algorithms: weak parameters, obsolete algorithms and implementation defects require their own treatment. An asymmetric wrapping key protecting symmetric data keys is included, even though the bulk encryption is not.

If the applicable population is non-empty, layer 1 has CCI 10,000 and largest share 100%. When no such assets remain, report **Not applicable – no remaining assets in the specified class**, with an asset count of zero; do not manufacture a CCI of zero. This baseline is excluded from the service maximum so that changes at the other layers remain visible. It reads 10,000 until the last such asset retires. No procurement action moves it; only migration does.

A hybrid containing a classical public-key component remains in the baseline as an inventory convention. This does **not** mean that breaking its classical component breaks a sound hybrid. The failure-impact determination and Effective Coverage must recognise the composition's actual security properties.

Layer 2 – Implementation lineage

Identify the implementation performing the relevant operation: for example, the deployed cryptographic library, provider module or firmware implementation. Retain product, version, build and deployment identifiers. Group versions into an implementation family only under a documented rule relevant to the failure mechanism. Different marketing names neither establish nor disprove independence.

Keep these executing families as separate primary units when their identities are distinct. A shared ancestor creates a reach relationship; it does not automatically merge them into one CCI unit. The relationship is supported only to the extent that relevant code or a relevant dependency remains present. A fork's age, separate maintenance team or different brand is not sufficient evidence of divergence. Conversely, historical derivation does not establish that every defect in the ancestor survives in the current fork.

Every disclosed ancestor becomes a reach node (C.3.2), and its edges follow the pedigree: nested chains for derivation, flat entries for a genuine merge. An edge exists only where the relevant code remains present in the descendant, and evidenced divergence – recorded commits or patches, reproducible source comparison, or corroborated engineering attestation covering the relevant primitives – severs it. Where two implementations resemble each other and the direction of derivation cannot be established, record them as variants and treat the shared membership as bounded. Similarity of vulnerability alone never creates an edge.

KyberSlash was a secret-dependent division timing flaw disclosed over late 2023 and early 2024, affecting multiple Kyber implementations including ones derived from or

incorporating the reference implementation. A related 2024 compiler-induced leakage issue showed a second failure route: constant-time source compiled into a binary with secret-dependent behaviour. PQClean, which influenced a significant body of post-quantum implementation work, was archived read-only on 4 August 2026. mlkem-native is used as the ML-KEM implementation source by liboqs and AWS-LC, and downstream software such as rustls may obtain ML-KEM through providers incorporating those implementations.

Layer 2 applies before migration begins: the classical libraries executing today, resolved to their ancestors, are the baseline against which post-quantum lineage is later compared. Heartbleed and the Debian defect are the same failure shape. Running two stacks during migration raises lineage diversity temporarily; a falling layer-2 index during migration is an artefact of two stacks, not a durable reduction, and is read with Effective Coverage.

The KyberSlash research demonstrates the significance of implementation-level timing defects across several Kyber implementations. It supports investigation of shared code paths; it does not justify inferring common ancestry from similar vulnerabilities alone. Source provenance, supplier engineering evidence and deployment confirmation are separate parts of that investigation.

Layer 3 – Trust root

Resolve certificate-bearing assets to the anchor actually relied upon by the relevant verifier. Record its accepted paths and trust policy. An issuer name alone is not a root determination, and a provider operating several roots is not necessarily one cryptographic root.

Two reach modes apply at layer 3. **Path invalidation** – a bridge or parent anchor reaches every asset whose accepted paths its compromise would all invalidate. **Compulsion** – a trust-list operator, bridge operator or legal authority that can revoke or compel mis-issuance across cryptographically independent anchors is one reach node under the compulsion mode, recorded separately from key compromise.

Apply failure-mode-specific reasoning to alternative paths. For an availability scenario involving removal of one root, another accepted path may preserve operation. For an integrity scenario, an alternative safe path does not remove forgery exposure if the verifier still accepts chains under the compromised root. A shared trust-list operator or policy authority may create a further dependency, but its powers and the assets subject to them must be established rather than inferred from a common jurisdiction label.

Layer 4 – Key custody platform

Use firmware family as the primary unit, supported by the supplier's identifier and relevant deployment evidence. Record original manufacturer and product identity

separately. A manufacturer-wide reach analysis is required where several families have a common manufacturer, but a common legal manufacturer alone is not proof of one firmware defect affecting all families. State whether the scenario concerns shared code, hardware, manufacturing, service withdrawal or another relevant mechanism.

Manufacturer reach at layer 4 is mandatory wherever a multi-family manufacturer is evidenced or disclosed. An undisclosed manufacturer follows the standard bound handling.

For cloud custody, resolve to the underlying module where disclosed. Where the provider is the only known boundary, Do not insert its name as a firmware family, code ancestor or hardware manufacturer. Retain the service boundary and the unresolved assignments, and compute bounds where the known population permits them.

Module validation and algorithm validation have different scopes. NIST's CMVP evaluates cryptographic modules against applicable module requirements; the CAVP evaluates algorithm implementations. Neither should be represented as a general certificate of supply-chain independence. A security policy may nevertheless provide useful provenance evidence.

Layer 5 – Protocol and negotiation

Measure the cryptographic outcome actually selected, not merely offered or configured. For key establishment, the outcome tuple consists of protocol, version and key-establishment group or algorithm. Preserve a function discriminator when other functions are assessed. For authentication, record the protocol/profile and selected authentication method rather than placing an authentication method in a key-establishment field.

Two paths negotiating the same tuple share the dependency. The quantum exposure of negotiated outcomes is not re-measured at layer 5, because layer 1 and Effective Coverage carry it and collapsing all classical tuples into one unit would reproduce layer 1's saturated reading. Layer 5's question is outcome concentration under the downgrade failure mode.

The primary tuple must be reproducibly normalised. A TLS cipher-suite label alone does not identify all relevant negotiated parameters. A payment message format or a 3-D Secure version alone is not a cryptographic outcome. Sector extensions identify additional observations needed for their protocols.

Where several outcomes occur, preserve their observation windows and affected asset or endpoint populations. A modal outcome must not conceal a fallback population. If one endpoint has a genuinely mixed policy that cannot be partitioned into disjoint configuration identities, use a documented composite outcome unit for the CCI and retain each constituent outcome in the exposure and coverage records. Transaction

frequencies may support a supplementary view; they do not silently replace asset-count weighting.

Layer 6 – Key generation

Identify the component, module or operational generation service that produced each persistent key. Custody after generation is not evidence of origin: an imported key can be held in one HSM and have been generated elsewhere. The generation-point identifier and the chosen grouping of operational replicas must be recorded and kept consistent between assessments. A supplier name alone is not a generation point.

Two failure ways cross vendors at layer 6. A predictable entropy source makes every key it seeds recoverable. A defective generation algorithm – structured primes, biased nonces or flawed derivation – makes keys recoverable from their public forms.

The generation point is the exclusive unit. The entropy-source design and the generation-algorithm design are reach nodes reaching every point that embodies them. Two points may share a design; one point may rely on several design components. Preserve these relationships without adding another vote to the asset denominator.

Certification-profile conformance feeds a design's upper bound only: known reach is the largest single conformant point and the bound spans all conformant points, until disclosure collapses the interval either way.

The ROCA disclosure came from one key-generation library from one semiconductor vendor and crossed trusted platform modules, authentication tokens, smartcards and a national identity scheme, each product with its own brand, supply chain and certification. The Debian OpenSSL advisory concerned one patched-out line reducing the seed to a process identifier, with every key generated on descendant systems for two years predictable. Both propagated through lineage: layer 2 asks whose executing code descends from a common ancestor, while layer 6 asks whose key generation shares a design. Neither implies that all key-generation defects produce the same reach or retrospective consequence. The affected versions, generation conditions and protected functions determine the scenario.

A.5 ROLES AND INDEPENDENCE

Role	Responsibility
Accountable executive	Approves the mandate, accepts material residual exposure and assigns remediation resources
Assessment owner – second line	Owns the method's application, challenges

Role	Responsibility
	scope and evidence, records determinations and presents findings
Service owner – first line	Confirms service boundaries, criticality facts, operational effects, tolerances and recovery assumptions
Cryptographic inventory owner – first line	Supplies the inventory, service mapping, source records and discovery-coverage statement
Architecture and engineering – first line	Establish executing components, accepted paths, custody and generation relationships; explain technical limitations
Third-party risk and procurement	Obtain supplier disclosures and maintain contractual or commercial follow-up, according to institutional responsibilities
Operational resilience	Maintains the service/tolerance mapping and integrates scenarios into resilience assessment
Independent reviewer or model validation	Reviews calculations and material assumptions independently of their development and execution
Internal audit – third line	Provides independent assurance over the design and operation of the assessment and its governance

Third-party risk is second line and owns the lineage-request process. Determinations are made by the assessment owner, not negotiated; first line supplies facts and may challenge on factual grounds; it has no veto on treatment.

The assessment owner must be in second line for a conformant CCF assessment. First line provides evidence and may challenge factual errors; it does not determine the favourable treatment of its own unsupported submission. Disagreements and unresolved challenges must be recorded and escalated through the institution's governance.

Independent validation is a function, not automatically a third-line role. Its organisational placement follows the institution's model-risk and assurance arrangements. Internal audit should not be made responsible for developing or operating the measurement it later audits.

A.6 EVIDENCE STANDARDS

Evidence is attached to a claim, not to a document as a whole. A runtime record may establish a negotiated group without establishing the ancestry of the implementation that selected it. A supplier may be authoritative about firmware lineage while having no observation of how the institution deployed the product.

Code	Evidence basis	Appropriate interpretation
E1	Runtime observed	Direct observation of behaviour within a stated deployment, sample and window
E2	Binary confirmed	Confirmation tied to the deployed artefact; ancestry additionally requires traceable source/build provenance
E3	Inventory declared	A configuration, SBOM, management-plane or comparable inventory assertion
E4	Vendor attested	An attributable supplier statement, including a relevant supplier-authored security policy
E5	Inferred	A reasoned conclusion from indirect records, with the inference documented

The admissibility rules are claim-specific:

Claim	Strongest admissible	Weakest admissible	Not admissible

Claim	Strongest admissible	Weakest admissible	Not admissible
Algorithm and parameter set	E1 runtime observed	E5 inferred	–
Implementation ancestry	E2 where source or build provenance, or reproducible binary-to-source provenance, establishes descent; otherwise E4 supplier attestation with recorded corroboration against a certification record (CMVP or ESV security policy) – the basis stays E4 and the corroboration is a separate attribute (S3)	E5 inferred	E1; runtime observation cannot establish descent. A plain binary signature match evidences presence, not descent, and is not E2 for this claim
Trust anchor	E2 chain built and verified	E3 inventory declared	–
Firmware family and origin	E4 supplier attested	E5 inferred	E3 alone; an inventory does not hold the fact
Negotiated outcome	E1 runtime observed	E3 configuration declared, recorded as capability-only	–
Key generation design	E4 supplier attested, or E5 chained through certification records	E5 inferred	E2; a binary does not reveal a hardware design

Where a claim's strongest admissible tier is E4, an E4 value is not a deficiency; that is the position at layers 4 and 6 for most institutions. Tier is recorded per value, not per document. Evidence tier never changes a computation; it changes the stated confidence, the challenge questions and whether a determination survives review.

The codes describe evidence basis. They are not a universal ranking in which a lower number always overrides a higher one. For each material claim, assess relevance, deployment match, recency, specificity, independence and contradictions. Corroborating a supplier statement with another supplier-authored document does not turn E4 into E3 or establish independent verification. Preserve the source basis and record corroboration separately.

Claim	**Evidence needed for the stated conclusion**	**Insufficient by itself**
Algorithm and parameters	An attributable observation, deployment artefact or inventory record with its limitations	A product's general support list as proof of deployment
Implementation ancestry	Source/build derivation, documented retained code or a specific engineering attestation	Runtime traffic; a product name; a binary string match without provenance
Trust anchor	Accepted-chain and trust-policy evidence for the relevant verifier	Issuer name or a root present in an unrelated trust store
Firmware family and manufacturer	Deployment identity linked to supplier or certification evidence	Reseller identity or an unverified family inferred from branding
Negotiated outcome	Runtime/endpoint observation for the claimed population and period	Configuration capability presented as an observed outcome
Key-generation design	Provenance linking the generation point to a design and relevant operating conditions	A generic RNG class, certification profile or current custody location

The Profile provides `conc:evidenceTier` as a component-level default and more specific `lineage:evidence` and `entropy:evidence` properties. These cannot encode every per-claim distinction. The assessment supplement must therefore retain material exceptions, evidence references and dates at claim level. A default must not override a more specific record.

An ungraded **supplied value** is provisionally E5 and identified as ungraded pending review. A missing value is unknown, not an inferred fact. Evidence grades do not weight the CCI or discount asset counts. New evidence can change an attribution or a reach relationship through a recorded determination, in which case the resulting calculation changes for that reason.

A.7 DETERMINATIONS AND UNCERTAINTY

A determination records the facts, treatment, reasoning, evidence, accountable decision-maker and challenge outcome for a material judgement. It is required where resolving a dependency, counting an asset or interpreting a failure requires more than a mechanical application of established rules. Use the Determination Log.

Four rules bind determinations. Resolve toward the failure mode – would one defect take both, not are these the same product. Absence is never a favourable value – an unassessed layer is recorded as not assessed, never as diversified. Non-disclosure widens a reach bound and an unresolved boundary is counted as one primary unit at its known population while uncertainty remains in reach. Record the determination, not just the outcome.

Resolve dependencies against the stated failure mechanism. Do not treat a common name as proof of common failure, or missing disclosure as proof of independence. Where supported alternatives remain possible, report bounds or a documented limitation rather than select the favourable result. Decisions must be reproducible from their records, including the basis for excluding a possible shared relationship.

Status applies to different dimensions; it is not one mutually exclusive enumeration covering all uncertainty:

Status	Meaning and reporting treatment
Not assessed	An applicable layer or output was not computed because the required input was unavailable
Not applicable	The layer or output has no applicable population under a documented scope rule; not the same as missing data
Excluded	A particular asset, layer or view is intentionally outside a specified calculation, with the reason retained. Examples: post-quantum material at layer 1; layer 1 in the reported maximum.

Status	Meaning and reporting treatment
Bounded	A numerical interval represents unresolved assignment or membership; identify the quantity and assumptions bounded. Examples: undisclosed lineage at layer 2; profile-only design evidence at layer 6. The interval runs from evidenced edges to the widest membership that cannot be excluded.
Unresolved boundary	A visible service or provider boundary conceals the identity needed for resolution. Example: undisclosed cloud custody across layers 2, 4 and 6. The boundary's identity is never substituted into sublayer fields.
Hybrid-unverified	A hybrid was identified but the evidence required by C.6 is insufficient; it receives no confirmed target-state credit
Unknown coverage	The completeness of the relevant inventory or path population cannot be established
Not assessable	A tolerance determination cannot be completed because required operational facts or thresholds are missing

A layer can have a computed CCI and bounded ancestry reach. A service can have an unqualified data-level claim over a stated subset while whole-estate discovery coverage remains unknown. Reports must identify these distinct limitations instead of compressing them into a single confidence label.

PART B – ASSESSMENT METHOD

Phases 0 to 2 are preparation, Phase 3 is the substantive work, Phases 4 and 5 are mechanical, and Phases 6 and 7 turn the assessment into a decision. A first cycle should expect Phase 3 to take more effort than all other phases combined.

The eight phases establish a traceable route from mandate to decision. They need not be implemented as eight separate projects. In a first cycle, data collection and supplier requests commonly overlap with resolution, but no reported calculation should lose the record of the scope and determinations on which it depends.

PHASE 0 – MANDATE AND SCOPE

Purpose. Establish the decision, accountable roles, perimeter and intended use of the assessment.

The assessment owner obtains the service list, applicable risk policies and existing tolerance statements. The charter names the accountable executive, entities and jurisdictions in scope, reporting audience, intended decisions, target data level, planned coverage analysis and resources. It also records the exact editions of the Universal Framework, Profile, Conformance Statement and sector extensions to be used.

Do not set CCF-3 as the target where suppliers have never been asked for lineage. Set CCF-2 and make the disclosure request a Phase 2 activity.

Set a feasible initial data target. A supplier-dependent evidence gap should become a planned workstream, not an unstated assumption that advanced inputs will be available. External reporting may require additional assurance and disclosure controls; those requirements are identified at the outset.

Common failures. Scoping to systems rather than services; setting a conformance target the data cannot support.

Outputs and exit criteria. An approved charter, a second-line assessment owner and an agreed decision use. Scoping only to convenient systems, without the services they support, is a scope defect to resolve before presenting a service-level result.

PHASE 1 – SERVICE SELECTION AND BOUNDARY

Purpose. Define what the service assessment includes and what target state and tolerance it will test.

Select services by criticality and by coverage of the cryptographic estate, not by data availability. A first cycle SHOULD run five to eight services; steady-state cycles run five to twelve. A well-scoped assessment of one critical service can be useful; a larger institution may require a wider programme.

Map each service to its applications, infrastructure, cryptographic endpoints, issued estates and external dependencies. Shared infrastructure is included in every service it supports. Determination example: one custody platform supports eleven services and an analyst proposes one eleventh each. Treat it as full exposure in every service. If the platform fails all eleven fail, and dividing reports each service at an eleventh of its actual exposure. Its assets are not divided by the number of supported services: each service remains exposed to the relevant failure. Within a service, duplicate observations of the same asset are removed under Annex 2.

Record the service's own tolerance or the inheritance determination in C.5.4. Define its cryptographic target state by function, including any distinct key-establishment and signature requirements. State the accepted algorithms, parameter sets, hybrid conditions, observation period and material fallback policies. A service may require both confidentiality and authentication protection; completing only one track must not be described as completing both.

Common failures. First line selecting the services; a boundary excluding shared infrastructure because it "belongs to platform engineering"; target state left implicit.

Outputs and exit criteria. Each service has a boundary, owner, asset-class scope, target-state definition and tolerance or explicit `NOT MAPPED` record. The assessment supplement identifies any service-specific facts that cannot be expressed unambiguously on a shared inventory component.

PHASE 2 – DATA REQUEST AND RECEIPT

Purpose. Obtain the evidence package and establish its completeness before interpreting the figures.

Issue the Data Request to the inventory owner. Request the CBOM, service and relationship mappings, coverage statement, evidence records and assessment-specific facts required for the intended outputs. The request distinguishes required applicable fields from optional corroboration and facts that do not apply to the submitted asset type.

Issue the Supplier Lineage Disclosure Request early for material unresolved components. A request unanswered at its deadline is `no-response` and a named non-response. Explicit refusal is `declined`. A response citing certification without addressing provenance is recorded `declined` with reason deflected, and counts as a received, unusable response in the disclosure register.

Issue lineage requests through third-party risk in Phase 2 even where the target is CCF-2, because supplier response cycles run to weeks and the response rate is itself a reported figure.

Challenge discovery coverage against an independently enumerated estate. A first-line claim of 92% coverage derived from certificate lifecycle tooling is recorded as unknown for classes the tooling cannot observe, with the observable perimeter stated. A coverage claim is as broad as its discovery method. A certificate tool that covers 92% of registered certificates has not demonstrated 92% coverage of embedded keys, firmware or issued devices. Preserve the scoped percentage and record the other classes as unknown. Where a class has an independently known population but no observations, its observed coverage can be zero; that is different from concluding that the class contains no assets.

Inspect the received package for secrets and access restrictions. The request is for metadata, public identifiers and evidence, not private keys or generation seeds. Protect detailed topology and supplier information even when it contains no key material.

Outputs and exit criteria. A registered input snapshot, a challenged or explicitly qualified coverage statement, evidence defaults and exceptions, and a disclosure-request register. Gaps are accepted and reported; they do not acquire favourable default values.

Phase structure. Phases 0 to 2 are preparation, Phase 3 is the substantive work, Phases 4 and 5 are mechanical, and Phases 6 and 7 turn the assessment into a decision. A first cycle should expect Phase 3 to take more effort than all other phases combined.

Phase 0 common failures. Scoping to systems rather than services; setting a conformance target the data cannot support. Where suppliers have never been asked for lineage, do not set CCF-3 as the first-cycle target; set CCF-2 and make the disclosure request a Phase 2 activity.

Phase 1 rules. A first cycle SHOULD run five to eight services; steady-state cycles run five to twelve. Select by criticality and coverage of the cryptographic estate, not by data availability. For shared infrastructure, count the asset in full in every service it supports. If one custody platform supports eleven services, allocating one eleventh to each understates every service's exposure because failure of the platform affects all eleven. Record `NOT MAPPED` where no tolerance and no inheritance exist. Common failures are first line selecting the services, excluding shared infrastructure because it belongs to platform engineering, and leaving target state implicit.

Phase 2 determinations. Issue lineage requests through third-party risk even where the target is CCF-2, because supplier response cycles run to weeks and the response rate is itself reported. A first-line claim of 92% coverage derived from certificate lifecycle tooling is not 92% estate coverage: record coverage as unknown for classes the tooling cannot observe and state the observable perimeter. Silence after the deadline is `no-response`; explicit refusal is `declined`; a deflected reply is `declined` with reason deflected and remains a received, unusable response in the disclosure register.

PHASE 3 – RESOLUTION

Purpose. Convert inventory records into primary assignments, reach relationships and explicit unresolved cases.

Apply Annex 2 to establish counted asset identities. Determine layer applicability and the primary unit for each applicable asset. The resolution ledger must retain the asset reference, service, layer, in-scope function, unit identifier, evidence and relevant determination. Record additional component relationships separately so that non-exclusive exposure is not lost.

Worked determinations.

1. **Parameter sets at layer 1.** Fact pattern: RSA-2048, RSA-4096 and P-256. Treatment: one dependency. Reasoning: the specified quantum failure mode is common.
2. **Symmetric material with asymmetric wrapping.** Fact pattern: AES-256 backup data keys wrapped by RSA-4096. Treatment: AES keys are excluded from layer 1; RSA wrapping keys are included and meet criticality test 1. Reasoning: the wrapping keys protect the symmetric estate.
3. **Convergent ancestry.** Fact pattern: two suppliers' libraries take ML-KEM from the same upstream. Treatment: two executing families remain two primary units, and the upstream is one reach node with evidenced reach over both. Reasoning: the index measures the executing estate and the upstream defect reaches both through shared code where it remains present.
4. **Divergent fork.** Fact pattern: six-year independent maintenance with documented divergence in the relevant primitives. Treatment: no edge for those primitives, at E4. Divergence claimed but not evidenced is treated as shared and logged for challenge.
5. **Non-disclosure disguised as disclosure.** Fact pattern: "industry-standard, FIPS-validated cryptographic implementations". Treatment: record `UNDISCLOSED`, disclosure status `declined` with reason deflected, and carry the unit in every non-excludable reach bound. Reasoning: validation confirms correct computation, not ancestry.

6. **Undisclosed ancestry.** The undisclosed unit stays in the index as itself; every reach node whose membership cannot be excluded carries it in its upper bound. The spread is the cost of non-disclosure in reach points.
7. **Intermediates under one root.** Four intermediates, three under one root: two dependencies, not four. Reasoning: compromise or deprecation of the shared root reaches all three intermediates at once.
8. **Cross-signed anchors.** Record both anchors, cross-signing and paths actually accepted. One dependency exists only where compromise, deprecation or removal of a common anchor invalidates both accepted paths; there are two where an independent valid path survives. Reasoning: cross-signing establishes a relationship between certification paths, not a common-mode dependency; the dependency is resolved against the failure mode and the paths the service relies on.
9. **Policy root versus cryptographic root.** Under key compromise the anchors stay distinct; under compulsion the authority is one reach node spanning every anchor it governs. The failure-impact determination selects the mode. eIDAS trust lists and federal bridge structures are the pattern.
10. **Rebadged modules.** Two suppliers, one OEM, one firmware family: one dependency. Reasoning: this is the case a vendor-layer substitutability test cannot see, and the reason layer 4 resolves to firmware family rather than brand.
11. **Cloud custody with undisclosed module.** The provider is the unresolved visible boundary and is counted as one primary unit at its known population. Lineage, firmware-family, manufacturer and generation fields never carry the provider's identity. The same boundary at layers 2, 4 and 6 is itself a finding.
12. **Capability versus outcome.** A hybrid-capable acquiring path negotiating classical on every connection is classical at layer 5 and contributes zero to Effective Coverage. Reasoning: the outcome on a path is set by the weakest accepted link, and recording capability would report a migration position the institution has not reached.
13. **Shared certification profile is not shared design.** Two platforms conforming to one RNG profile with no design disclosed: for that bounded design node the known reach is the largest conformant point and the bound spans both. A profile specifies properties a design must meet, not the design.

An undisclosed ancestor normally changes reach uncertainty while leaving known primary assignments unchanged. If disclosure establishes that a primary assignment itself was wrong – for example, two purported families are actually one family – correct the CCI as well and classify that movement as re-attribution. "Disclosure never changes the index" is not an exemption from correcting the inventory.

Outputs and exit criteria. Primary-unit assignments, a reach/relationship ledger and layer applicability are recorded. Every asset is resolved or recorded as unresolvable, and unresolved cases remain visible. Every determination is recorded with reasoning

sufficient to reproduce its treatment. Architecture confirms material technical resolutions, including layers 2, 4 and 6. Common failures are resolution to product name at layer 4, marketing language accepted as lineage disclosure, and ambiguity resolved to the favourable reading.

PHASE 4 – COMPUTATION

Purpose. Compute the structural measures over the resolved populations.

For each service and layer, record its applicable asset count and compute the CCI, largest share, number of primary units and criticality views under Part C. Compute named-domain reach and its bounds. Use unrounded values for calculations and round only the presentation.

A missing layer is not a zero in the service maximum. A bounded assignment is not a known unit selected for convenience. Empty flagged subsets have no flagged-only CCI; label them as no qualifying assets, rather than diversified.

Outputs and exit criteria. Reproducible calculation records for every computable cell, and explicit statuses for the others. Independent arithmetic checks should be performed before results are interpreted.

Phase 4 required outputs. Every cell carries the index, largest share, flagged-only index and reach table with known reach and upper bound per node. Every non-computable layer is not assessed, never diversified. Common failures are averaging across layers and omitting the largest share.

PHASE 5 – COVERAGE AND COUNTERPARTIES

Purpose. Establish the extent to which service assets meet the target state across their relevant paths and identify the parties governing the shortfall.

Read actual outcomes, own-state evidence, counterparty state and required hop information together. Apply C.6 to all relevant paths for an asset, including secondary operating paths that expose the same protected function. Counterparty roadmap dates support planning; they are not evidence that a deployed path already operates at target state.

Assign counterparty classes by the practical route through which the institution can secure change. A supplier may be contractual for one service and an effectively non-substitutable blocking dependency for another. Document the basis. Count distinct counterparties within a stated identifier and legal-group policy, not repeated occurrences of the same counterparty on multiple paths. Keep population measures separate from counts of individual organisations.

Outputs and exit criteria. EC or a qualified coverage interval per service/function, the EC data marker, counterparty counts by class, population denominators and the evidence gaps preventing stronger conclusions. Where the necessary population is unknown, report EC as not computable rather than as zero.

Phase 5 determinations. An asset on three paths, two at target and one terminating at an unmigrated counterparty, is not covered because coverage is a path property. A downstream provider that is contractually obliged but unmigrated is contractual, not blocking, because classification follows the remediation route. Common failure: counting a retail customer estate individually instead of as a population.

PHASE 6 – INTERPRETATION AND CHALLENGE

Purpose. Determine which structures and evidence gaps require action.

Apply C.4's bands as descriptive indicators. Test material failure domains under C.5 using service-owner recovery facts and consequence thresholds. Identify the scope of each result: evidenced domain, upper-bound scenario or incomplete operational determination. A high CCI does not by itself establish a tolerance breach, and a low one does not excuse testing a critical single dependency.

Challenge primary assignments, retained-code assumptions, exclusions from reach bounds, criticality judgements and the coverage denominator. Review proposed remediation against the actual dependency it changes. A supplier change can improve the commercial position without reducing a shared code failure domain; the assessment should show both effects rather than confuse them.

Outputs and exit criteria. Challenged findings, tolerance determinations, material evidence limitations and proposed responses. Material unresolved challenges accompany the result to the decision-maker.

Phase 6 challenge. A tolerance result asserted without a recorded determination is reversed in challenge. A layer at 5,200 with a largest share of 52% on a service with a 72-hour tolerance differs materially from the same index on a two-hour service. Common failures are reporting band positions as findings without materiality and skipping challenge on convenient determinations.

PHASE 7 – REPORTING, REMEDIATION AND RE-ASSESSMENT

Purpose. Convert the assessment into accountable decisions and maintain a comparable record over time.

Produce the service-level heat map and supporting findings. Apply the M3 gate before publishing an institution-level figure. Report owners, actions, dates, dependencies and

any risk acceptance alongside the figures. Use the Board Report Template for the decision summary; it does not replace the underlying matrix or determination records.

Assess at least annually for a stable estate and at least semi-annually during active migration, with event-driven review where a material implementation, trust, custody, generation or counterparty change occurs. More frequent automated collection may support this cadence without replacing periodic judgement and challenge. Decompose changes under C.8 before claiming remediation progress.

Outputs and exit criteria. A delivered decision pack, approved action or risk-acceptance records, named owners and the next review date. Record the model-classification determination under G.6 and maintain the supporting documentation irrespective of that classification.

Phase 7 rule. Sequence remediation by effect, not ease, and distinguish E.2's false moves. Common failure: delivering the institution-level figure without its companions.

PART C – METRIC SPECIFICATION

C.1 POPULATION, SHARE AND WEIGHTING

A counted asset is a persistent key, certificate, cryptographic endpoint/configuration or issued instance identified under Annex 2. Algorithm descriptions, software-library catalogue entries and path-state records support resolution; they are not additional counted assets merely because they are separate CBOM components.

The normative weighting is asset count. A cohort record may represent a known number of distinct issued assets, in which case its weight is that cardinality, not one. Transaction, financial-value or operation-weighted results may be useful supplements, but must be labelled separately and must not replace the normative count-weighted result.

Asset count is a pragmatic basis for consistent calculation, not an assertion that every asset has equal business value. A small number of master or settlement keys may be more consequential than a large population of routine certificates. C.2's criticality views and C.5's service-effect determination address that limitation; neither eliminates the need to state it.

Across the engagements informing the framework, asset count is the ceiling of what institutions can produce; transaction-volume weighting would produce a specification no institution could compute. The declared bias is that a settlement signing key and a thousand development certificates count as one and a thousand. C.2 is the compensating view.

Layer	Exclusive unit for the index	Standard reach nodes
1 Algorithm	Algorithm class under the assessed failure mode	–
2 Lineage	Executing implementation family	Ancestor codebases

Layer	Exclusive unit for the index	Standard reach nodes
3 Trust root	Terminating anchor of the primary accepted path	Bridge and cross-signing anchors; policy authorities
4 Custody	Firmware family	Manufacturer, where one manufacturer spans multiple families; evidenced shared hardware or SDK domains
5 Protocol	Negotiated outcome tuple	Downgrade-target tuples, where a forcing mechanism is evidenced
6 Key generation	Generation point	Entropy source designs; generation algorithm designs

Manufacturer reach at layer 4 is mandatory wherever a multi-family manufacturer is evidenced or disclosed. An undisclosed manufacturer follows the standard bound handling.

C.1.1 Layer applicability and denominators

For service `s` and layer `L`, let `A(s,L)` be the deduplicated population of assets to which that layer applies, within the stated assessment perimeter. Let `N(s,L)` be its asset count, including declared cohort weights. A layer's shares use that denominator, not automatically the entire service inventory.

Report `N(s,L)`, the counted asset classes, exclusions and unresolved counts with each cell. A certificate-bearing path and a symmetric key need not have the same layer applicability. Do not invent trust anchors for assets with no applicable certification relationship or generation points for assets whose relevant function does not involve generated key material.

Where discovery covers only part of the applicable estate, the calculation is a result over that covered population. It must not be presented as a whole-estate result without a supported coverage statement. If the number of missing assets is unknown, a narrow interval over the observed inventory does not bound the unknown whole estate.

If `N(s,L) = 0` because no assets apply, the CCI and largest share are **Not applicable**, not zero. If the population cannot be established, record the calculation as **Not assessed** or explicitly limited to a known subset. These statuses must be distinguishable in the reporting matrix.

C.1.2 Exclusive primary assignment

Each applicable asset contributes once to one primary unit in a cell. The primary assignment follows the relevant operational function and accepted execution path, with the rule recorded in the resolution ledger. Additional relevant implementations, fallback paths, trust dependencies and common designs are retained in reach and coverage records.

Where an operation necessarily uses several simultaneous implementations – for example, a classical/PQC hybrid and its combiner – the primary unit is a **composite execution unit** defined by the role-labelled tuple of those implementation families. Count the asset once against that tuple. Record each constituent implementation and shared ancestor as a reach dependency, with the failure semantics of the composition. Do not assign a full asset vote to both components in the same CCI denominator.

This composite-unit convention also applies where a genuinely inseparable primary configuration contains multiple relevant outcomes or generation components. It must not be used to hide avoidable uncertainty: first resolve disjoint endpoint/configuration populations where the evidence supports doing so. A lower CCI created by several composite tuples is not proof of independent security, because a component common to all tuples may still have 100% reach.

C.1.3 Unresolved primary identities

An unresolved boundary is counted as one primary unit at its known population. Its identity is never substituted into sublayer fields, and its uncertainty is carried in reach as a bounded node. Do not create unconstrained lower and upper CCI bounds by treating every unknown asset as its own unit or by joining every unknown asset to the largest known unit. If the population itself cannot be established, publish the known-subset result and the limitation.

C.2 CRITICALITY AND THE FLAGGED-ONLY VIEW

An asset qualifies for the criticality view if any of the following service-specific tests is true:

1. It certifies, derives or wraps other keys, including a relevant root/intermediate authority, master key, derivation master or key-encryption key.
2. It authorises an irreversible action or supports a finality or non-repudiation guarantee.
3. It cannot be revoked or replaced within the service's impact tolerance.
4. Its operational or verification validity extends beyond the institution's applicable migration completion date. Where an applicable external requirement supplies the horizon, record its scope and status; a draft date is not a binding requirement.

5. Its failure leaves no operational equivalent for the relevant service function.

The service or key owner supplies the facts. The assessment owner records the determination and challenges unsupported conclusions. A certificate, signing key or master key should be flagged for the test it actually meets; association with a payment service alone is not enough to prove every test.

The Profile carries five corresponding boolean properties. An omitted boolean is **unknown**, not false. If any known test is true, the asset is qualifying. If all applicable tests are established false, it is non-qualifying. Otherwise its criticality is unresolved. Record the known qualifying count and the unresolved count; do not present a complete flagged-only view where unassessed assets could materially change it.

A shared asset may have different criticality at different services because their tolerances and horizons differ. Store the asset-by-service determinations in the supplement when one component-level set of booleans cannot represent them faithfully.

A primary unit is described as flagged when it carries at least one qualifying asset. However, the flagged-only CCI is computed **only over qualifying assets**, not over all assets attached to flagged units. Recompute the shares, largest share and reach on that subset. Publish its asset count and the count of flagged units. If there are no qualifying assets, report **No qualifying assets** rather than a numerical CCI. If flags were not assessed, report **Not assessed**.

C.3 INDEX AND FAILURE-DOMAIN REACH

The CCI and reach answer different questions and are reported together wherever the relationship data permits. A computed CCI remains useful when ancestry is unknown, provided that the missing reach evidence is explicit. It must not be interpreted as evidence of upstream independence.

Every cell reports both the CCI and reach; a cell reported with one and not the other is incomplete. The layer-2 index computes at CCF-1 with no supplier data. Reach consumes lineage and design evidence as it arrives at CCF-3. Evidence that corrects a primary unit's identity moves the index and is a C.8 measurement effect.

C.3.1 Cryptographic Concentration Index

For a non-empty cell, let `n_d` be the asset count assigned to primary unit `d`, and `N` the applicable population. Define `s_d = n_d / N`, with shares between zero and one and their sum equal to one.

$CCI(L) = \Sigma\ s(d)^2$ for d = 1 ... n

For `n` equally sized primary units, the result is `10,000/n`; for one unit it is 10,000. The scale is an HHI-style concentration convention. It does not import competition-law thresholds or establish empirically calibrated security limits.

Report the largest primary-unit share `s_max`, its identifier and the number of units. Retain the unrounded shares and CCI in the calculation record. Presentation rounding should be consistent; examples in this suite round CCI to the nearest integer and shares to one decimal place.

A service's headline structural result is the maximum of its computable layer-2 to layer-6 indices, subject to C.7's incomplete-input rules. Layer 1 is shown separately. Averaging layers would allow a less concentrated layer to offset another layer's common dependency even though the two are different failure mechanisms.

Learning an ancestor normally changes the reach record, not a correctly resolved primary assignment. Correcting an executing-family identity, service scope or asset population can change the CCI. The explanation of the change is therefore part of the result, not an optional narrative.

C.3.2 Failure-domain reach and its bounds

Every reach edge carries its failure mode, evidence tier and one of two states: *evidenced* or *bounded*. A bounded node reports an interval from known reach to the upper bound, and the spread is the cost of non-disclosure, stated in reach points.

For a specified node `v`, failure mode `f` and cell, reach is the fraction of the applicable asset population exposed to that node under that mode. Each asset is counted once within a node's reach even if several paths connect it to the node. The same asset can be in several different nodes' reach sets; their percentages must not be added to estimate a total exposure.

For a named node with confirmed affected-asset set `K` and additional possible set `U`, report:

R_known = |K| / N, R_upper = |K ∪ U| / N

Cardinalities include declared cohort weights. Every included or excluded relationship must have a basis. The upper bound covers memberships that cannot be excluded under the documented candidate relationships; it is not a probability that those memberships exist. State the spread in **percentage points**, not as a change in the CCI.

For example, an ancestor established for 2,400 of 4,200 assets, with a further 900 assets not excludable, has 57.1% known reach and a 78.6% upper bound. The 21.4-percentage-point spread is an evidence limitation concerning those 900 assets. The other 900 assets require an explicit basis for exclusion from this ancestor's bound.

Evidenced means the membership is established. **Bounded** means it cannot be excluded: undisclosed lineage, or a design evidenced only by certification-profile conformance.

For that design node the known reach is the largest conformant point, because a design at least spans its own point: a floor by construction, not evidence that two points share a design. Where analysis is limited to known ancestor names, record the residual population whose commonality is unidentified rather than treating it as independent.

For the **largest failure-domain reach** companion, consider both primary units as standalone domains and additional upstream domains within layers 2 to 6. Layer 1's saturated quantum baseline is reported separately, as it is for the CCI roll-up. It cannot be smaller than the largest applicable primary-unit share merely because the table of named upstream nodes contains only smaller nodes. Identify whether its endpoint is a standalone unit, a named upstream node or a bounded design. Every cell reports its largest failure domain – the larger of its largest unit and its largest reach upper bound – with that kind label. A cell with no shared or bounded upstream reports its largest unit, so no cell shows a dash.

The diversity ladder has three rungs only: supplier/brand count, primary-unit count, and the reach view. There is no fourth rung. It does not define a single count of "effective independent domains": overlapping reach sets do not form a partition. When criticality is assessed, calculate reach over qualifying assets separately as well.

C.4 BANDS AND LARGEST-SHARE INTERPRETATION

The bands derive from the largest-share bound rather than from merger-review thresholds, which encode market power and say nothing about defect propagation. The bands are descriptive design conventions. Their breakpoints reflect squared-share lower bounds, not loss history, a forecast of cryptographic failure or a validated acceptance threshold.

CCI range	Band	Reading
Below 2,500	Diversified	No single dependency reaches half
2,500 to below 4,900	Concentrated	One dependency can take half
4,900 to below 8,100	Highly concentrated	One dependency can take the majority

CCI range	Band	Reading
8,100 to 10,000	Single-source-equivalent	Treat as common-mode

Exact boundary values belong to the higher band. "Diversified" means dispersed across the **primary units measured in that cell**; it does not establish independent ancestry or acceptable residual risk. "Single-source-equivalent" is a reporting band, not a statement that there is literally one supplier.

If the largest share is m, the CCI is at least $10{,}000 \times m^2$. With a fixed number $n > 1$ of units, the equal distribution of the remainder gives the sharper minimum $10{,}000 \times [m^2 + (1-m)^2/(n-1)]$. Thus a 50%, 70% or 90% share requires a CCI of at least 2,500, 4,900 or 8,100 respectively. The converse does not hold: a CCI above one of those boundaries does not prove that the largest share exceeds the corresponding percentage.

Two equal units produce CCI 5,000 and largest share 50%, not 70%. Report s_1 in every cell. Report the measured largest share, reach and criticality alongside the band. Where an interval crosses a band boundary, show both the interval and the bands it spans rather than selecting the favourable band.

C.5 FAILURE-IMPACT DETERMINATION AND TOLERANCE TEST

A structural share cannot establish how much of a service fails or how long recovery takes. The failure-impact determination supplies that transmission step. It is a scenario assessment using operational facts, not a mechanical conversion of CCI into downtime.

C.5.1 Required record and selection of domains

For each tested domain, record five things: 1 failure scenario under C.5.2; 2 affected assets and functions; 3 service effect – halts, degrades, or continues with a compromised property; 4 recoverability – time to detect, decide and remediate, estimated by the service owner and challenged; 5 result – Fail, Pass, or Not assessable, never Pass by default. Recovery includes detection, decision, emergency change, re-keying or replacement, redistribution and the restoration of dependent parties where relevant. Distinguish estimated time from tested capability.

At a minimum, examine the largest primary unit; the additional reach domain with the largest upper bound where different; the largest flagged primary and reach domains where applicable; and material domains carrying test-2 or test-5 assets. The selection is not limited to four records: several critical domains may need assessment, and duplicate selections need not be repeated.

A tested domain is **Fail** where the scenario exceeds a defined tolerance or produces an intolerable irrecoverable consequence. It is **Pass** where the supplied operational

evidence supports remaining within all applicable limits. It is **Not assessable** where a necessary fact or threshold is absent. The cell result is Fail if any required test fails; otherwise Not assessable if any required test is incomplete; otherwise Pass. Preserve the individual results behind that summary.

For bounded reach, assess the upper-bound scenario and label the conclusion accordingly: for example, **Fail – upper-bound scenario; shared membership unconfirmed**. This means resilience has not been demonstrated against that permitted scenario. It is not evidence that all possible members actually share the dependency, that an incident has occurred, or that the bound is a likelihood estimate. Operational facts still have to support the result; a large bound alone does not supply a recovery time.

C.5.2 Layer-specific scenario semantics

Layer	Scenario to specify	Consequences and response to examine
1	Break of the specified public-key class	Confidentiality or forgery consequences; whether a sound composition preserves the relevant property; replacement of vulnerable functions
2	Exploitable defect in a retained implementation path	Affected versions and functions; patching, rotation, deployment coordination and any preceding exposure
3	Anchor compromise, trust withdrawal or common policy action	Accepted forged chains, surviving valid paths, revocation, reissuance and trust-store redistribution
4	Firmware/module defect or an applicable withdrawal/recall	Availability and security effect; key portability, substitution, validated operating constraints and recovery capacity
5	Failure or forced selection of a cryptographic outcome	Exposed paths, ability to disable the outcome, interoperability loss and

Layer	Scenario to specify	Consequences and response to examine
		alternate routes
6	Defect in the source or generation design	Which generated keys are affected, replacement of the generator, re-keying, distribution and residual exposure of prior data or signatures

The scenario must not assume that every defect in a layer compromises every asset. Version, function, retained code, composition and operating conditions qualify the affected set. A hybrid may survive a break of one mathematical component and still be vulnerable to a shared implementation defect that compromises both components or the endpoint.

C.5.3 Confidentiality and integrity thresholds

Impact tolerances may include time and other service-impact measures. A time-only disruption limit is insufficient for a confidentiality failure that leaves the service available, or a forged authorisation whose consequence cannot be reversed by restoring availability.

Where relevant, the institution must define a confidentiality or integrity consequence threshold and test it alongside the operational tolerance. The threshold might identify prohibited exposure categories, an unacceptable class of forged instruction or another measurable harm. The sector templates illustrate how to record it; they do not prescribe a legal or universal zero-loss threshold.

An undefined consequence threshold is never a pass; its absence is a finding for the operational resilience function. Missing thresholds produce Not assessable results and a governance finding. Remediation after disclosure cannot be assumed to reverse historical confidentiality loss or restore the authenticity of every earlier record. Assess preservation, revocation, timestamping or reissuance controls on their actual properties rather than on a generic "re-keyed" label.

C.5.4 Tolerance inheritance and unmapped services

Rule 1. Use a service's own approved tolerance where available. **Rule 2.** If a supporting service lacks one but supports identified important business services, use the most stringent relevant inherited tolerance and document the relationship. An inherited tolerance is not permission to redefine the parent service's scope.

Rule 3. Where the service supports none and no applicable mapping can be established, record `NOT MAPPED`, reported as not applicable, never as passed. Do not record Pass. Report both the number of unmapped services and the number of Not assessable cells: the first concerns service governance; the second may arise even where a tolerance exists but operational facts or consequence thresholds are missing.

C.6 EFFECTIVE COVERAGE

Effective Coverage is the share of the service's in-scope persistent cryptographic assets for which all relevant operating paths and required cryptographic functions meet the service's defined target state. It is an asset-weighted measure, not a percentage of connections, counterparties or products that advertise support.

C.6.1 Calculation and evidence

Let `A_EC` be the identified asset population to which the target applies. For asset `a`, let `P(a)` be its relevant paths, and let `q(a,p)` be one when the required own-state, counterparty and hop conditions are established at target and zero when an applicable condition is established below target. For complete binary observations:

$$EC = 100 \times \Sigma\, w(a) \cdot \min \text{ over } p \text{ in } P(a) \text{ of } q(a,p) \,/\, \Sigma\, w(a)$$

`w_a` is one for an individual asset and the stated cardinality for a homogeneous, deduplicated cohort. Cohort members may share a weight only where they share the conditions relevant to the coverage determination. The calculation record must retain the asset denominator and its mapping to paths, functions and observations.

A path may include several cryptographic hops, such as termination and re-establishment at an intermediary. Establish the relevant protections on those hops and any required storage or authentication boundary. Do not multiply several copies of one observation or assume that protection on an outer tunnel proves every inner processing boundary meets the target. Equally, do not invent an exposed intermediate hop when encryption genuinely remains end to end; define the architecture being assessed.

Where one asset uses several relevant paths, all must satisfy the target for that asset to count as covered. A primary path at target does not remove exposure on a permitted fallback path. An asset with no path mapping is a mapping gap, not an automatically covered asset. For a wholly internal function, assess the relevant internal path or own-state conditions; external counterparty counts are zero and external counterparty fields are not fabricated. Where the function has no network path, represent its local protection conditions as a single assessment entry in `P(a)`; do not evaluate a minimum over an empty set or invent an external connection.

Unknown states must remain distinguishable from observed failures. On a known complete asset/path population, a **confirmed coverage lower bound** can count unknown

conditions as receiving no confirmed credit, while a corresponding upper bound allows unresolved conditions to be at target unless a known failing condition prevents it. Label that interval explicitly. Where the applicable population or its required path relationships cannot be established, report EC as not computable for that scope; a percentage over a selected known subset is a separate qualified result.

C.6.2 Hybrid acceptance

A hybrid can count at target only where the service's target definition accepts it and evidence establishes all three conditions:

1. Its post-quantum component and parameter set meet the applicable approved standard and deployment requirements. Key-establishment and signature standards serve different functions; FIPS 203 specifies ML-KEM, while FIPS 204 and FIPS 205 specify signature schemes.
2. The composition preserves the required security property when either component remains secure under the stated threat assumptions. For key establishment, merely concatenating two secrets is not proof of a robust combiner. For a dual-signature target requiring both signatures, both must be verified with the required binding and policy; accepting either signature is a different construction.
3. The operational path enforces the target and does not silently fall back to a below-target mode. Record observed and permitted fallback behaviour, scope and controls.

Soundness changes what a single break costs and changes no concentration figure.

A construction whose combiner or fallback properties cannot be established is `hybrid-unverified` and earns no confirmed target-state credit. A construction that is not accepted by the target does not become acceptable merely because its algorithm names include a post-quantum component.

The classical component remains in layer 1's inventory baseline. Layer 2 counts the composite asset once and preserves its constituent dependencies in reach. Hybrid acceptance changes the coverage and consequence assessment; it does not assert implementation independence or eliminate other layers' failure modes.

C.6.3 Populations and joint state

For populations, record the denominator, measurement date, segmentation and basis of `populationTargetPct`. A count of enrolled devices is not automatically the denominator for transaction-weighted readiness. The choice must match the count-weighted asset population being assessed.

Do not take the minimum of unrelated marginal population percentages as an exact joint coverage figure. If 80% of one cohort meets one required condition and 70% meets another, without evidence identifying the overlap, the share meeting both can be

between 50% and 70%. The minimum is an upper bound, not an observed intersection. For `k` required conditions measured on the same population, with fractions `p_j` and no additional dependence information, the joint fraction is bounded by `max(0, sum(p_j) - (k-1))` and `min(p_j)`.

Use joint observations, disjoint segments or an evidenced nesting relationship for an exact result. Otherwise carry the joint-state bounds through the asset-weighted calculation. Do not assume statistical independence unless an explicitly separate sensitivity analysis justifies it. The minimum-across-paths rule is exact for each identified asset's known binary states; it does not supply missing population overlap data.

C.6.4 Counterparty classification and markers

Class	Operational criterion	Typical response
Blocking	The required change cannot reasonably be compelled or substituted by the institution within the relevant programme	Industry coordination, escalation and contingency planning
Contractual	Contract, onboarding, certification or service-management rights provide a practical route to require change	Supplier/client remediation programme and enforcement of agreed obligations
Peer	Change requires coordination between parties with broadly independent authority	Bilateral planning and interoperability testing
Population	A distributed estate is managed through a common campaign rather than individual negotiations	Segmented rollout, refresh, certification and measured adoption

On the signature track, the certificate authorities in which a service's accepted chains terminate are blocking counterparties, except an institution-operated CA, which is an internal dependency. Their post-quantum issuance roadmap supplies `counterpartyTargetState`, and a roadmap date is not a deployed state. Where the institution runs the PQC Migration Framework, that roadmap arrives through its supplier questionnaire. Population class reports as a percentage. Coverage is the actionable

metric because it responds to sequencing decisions the institution controls, while concentration is largely inherited.

These are remediation classifications, not legal conclusions about particular organisations. A service's use of a public CA does not automatically make every CA blocking; an institution-operated CA is an internal dependency. An issuance roadmap is planning evidence, not a boolean assertion that a path has reached the signature target.

The Conformance Statement defines separate **EC-0, EC-1 and EC-2** markers for coverage-data sufficiency. They are independent of CCF-1/2/3. Layer-5 observation may be sufficient to compute outcome concentration while counterparties, target definitions or full path mappings are still insufficient for EC.

C.7 ROLL-UP AND REPORTING

C.7.1 Primary output

The primary output is a service-by-layer matrix. Each computable cell reports its asset denominator, CCI or bounds, largest primary share, criticality results, largest failure-domain reach and its basis, and tolerance conclusion. Show not-assessed and not-applicable cells distinctly. EC, counterparty classifications, discovery coverage and data-level claims are reported per service alongside the matrix.

Every quoted figure identifies its assessment date, perimeter, method/Profile revision, data sufficiency and relevant evidence limitation. The evidence distribution is reported by claim class or output where a single asset-level distribution would conceal material differences. A result over an incompletely observed estate must retain that qualification when copied into a board paper.

C.7.2 Service and institution figures

One figure is published because boards and supervisors will use one number, and an unpublished figure is replaced by one someone else derives from the service results.

A service's maximum is taken over layers 2 to 6. If one or more applicable layers are not assessed, call it a **computed-layer maximum**, state which layers are missing and do not imply it bounds the missing layers. Not-applicable layers are excluded with their scope rationale. A service with no computed applicable layer has no numerical maximum.

For services with supported weights and maxima, the institution-level indicator is:

$$CCI_{institution} = \Sigma\, N(s)\, M(s) \,/\, \Sigma\, N(s)$$

where `M_s` is the service's maximum or computed-layer maximum, and `N_s` is its declared deduplicated service asset count. Shared assets appear in every supported

service, so this is a **service-exposure-weighted mean**, not a deduplicated enterprise-asset mean. Layer-specific populations and the service weighting population must be retained separately.

If contributing results are intervals, propagate them without collapsing to a preferred endpoint, and identify conservative outer bounds where joint assumptions prevent asserting attainability. Unknown service weights, materially incomparable perimeters or uncomputed service maxima require a qualified subset roll-up or withholding, not invented weights. Report the number and identity of excluded services.

Withhold the institution-level indicator below assessment maturity M3. The gate applies to the assessment as a whole. M3 does not remove other input-quality limitations; a governed assessment can still need to withhold a number whose population is not defensible. Publish the service matrix and findings regardless of whether the institution-level indicator is available.

C.7.3 Mandatory companions and limitations

The institution-level indicator, where published, must be accompanied by:

1. The number of services at or above CCI 8,100, distinguishing confirmed from possible threshold crossings where bounds or missing layers prevent a single count.
2. The lowest EC across the assessed services, or an explicit statement that the minimum cannot be established because one or more services lack computable coverage. A minimum over known values alone must be labelled as such.
3. The number of services recorded `NOT MAPPED`.
4. The supplier disclosure response rate where requests were issued, with numerator, denominator, as-of date and pending requests.
5. The largest failure-domain reach across the assessed services at layers 2 to 6, including standalone primary domains and upstream domains, with its node or bounded design, bound and affected service.

Report the Not assessable cell count and other material data gaps with these companions. A companion not computable is reported as not computable, not omitted or replaced by zero.

Three limitations accompany every institution-level presentation. The mean can conceal a highly exposed service. Improving discovery can raise the index, and that is a sign the assessment is working. State this every time the figure is reported. Discovery and re-attribution can move the result in either direction without a deployment change. Asset-count weighting can understate the importance of low-count critical material, so the flagged-only and consequence views remain necessary. Cross-institution comparisons additionally require comparable service mappings, counting rules, layer scopes and evidence perimeters.

C.8 MOVEMENT BETWEEN CYCLES

Changes in a reported figure can result from changes to the estate, changes to measurement, or both. A declining CCI does not by itself establish remediation, and an increasing CCI does not by itself establish deterioration. The assessment must explain the change on a consistent basis.

Cause	Meaning	Class
Deployed	An asset entered service after the prior assessment	Estate
Retired	An asset left service	Estate
Re-platformed	An evidenced deployment change moved an asset's primary dependency	Estate
Discovered	An asset existed previously but was missing from the earlier inventory	Measurement
Descoped	An asset or service boundary left the measured perimeter without leaving operation	Measurement
Re-attributed	Improved evidence corrected an assignment without a corresponding deployment change	Measurement

Record causes at the level of the changed fact. One asset may have both a discovery correction and a subsequent deployment change; do not force an unsupported single cause when several events are known. Preserve the evidence and event dates needed to reconstruct the comparison.

For the index, use the suite's fixed decomposition convention: construct the current-cycle counterfactual with measurement changes reverted to their prior-cycle state where that state is supportable. **Estate movement** is counterfactual minus prior CCI; **measurement effect** is current CCI minus counterfactual. The terms sum exactly to the headline change before rounding. The decomposition is an accounting convention for a nonlinear measure, not a unique causal attribution independent of ordering.

For example, six assets on A and four on B produce CCI 5,200. A later count of eight on A and four on B produces 5,556. If one A asset was deployed and the other was discovered, the counterfactual is seven on A and four on B, or 5,372. The change is therefore approximately +172 real movement and +184 measurement effect, totalling +356. Neither component should be renamed a loss or probability change.

Report the common-estate share for **each** cycle: matched asset cardinality divided by that cycle's cardinality, using Annex 2's identity rules. When less than half of either population is common, mark the estate comparison indicative. If prior states cannot be reconstructed, report the decomposition as not estimable and explain why rather than assigning the residual to remediation.

For reach, use prior evidence on the current population as the counterfactual: population effect first, evidence effect second. Apply the split to the known and upper endpoints separately, retaining node identity and failure mode. Population effect is not automatically real remediation; it may contain changed discovery or scope and must be read with the cause records. A smaller upper bound supported by new disclosure is evidence improvement, not proof that the physical estate changed.

For institution-level movement, recompute the prior, counterfactual and current roll-ups, including changes in weights and service inclusion. Summing rounded cell deltas or holding changed weights fixed would not reproduce the headline movement. Remediation credit requires an identified beneficial estate change or a separately described improvement in evidence; direction of the index alone is insufficient.

C.8 reporting uses three rules: 1 never report a bare delta; 2 separate estate movement from measurement effect under the fixed counterfactual convention; 3 report common-estate share for each cycle and qualify comparisons where matching is weak.

Headline	Real movement	Measurement effect	Read as
Down	Down	Near zero	Remediation. Credit it.
Down	Near zero	Down	Composition or scope change, not remediation.
Up	Near zero	Up	Discovery working. Say so, per C.7.
Up	Up	Any	Concentration genuinely rising. Escalate per E.3.

Movement test vectors join the canonical set with the reference-implementation update in the v1.0 final package.

PART D – INTERPRETATION, CHALLENGE AND MATERIALITY

D.1 READING THE LAYERS TOGETHER

A high layer-2 CCI indicates concentration in primary executing families. A high ancestor reach indicates common code exposure across those families. Either may matter, but they are not interchangeable findings. The same distinction applies to generation points and the designs behind them at layer 6. A service with several generation points can still have one design domain; a large point need not establish that a design is shared elsewhere.

At layer 3, distinguish the distribution of accepted roots from the concentration of critical signing and certification functions. At layer 4, distinguish firmware-family concentration from the consequences of a manufacturer-wide event. At layer 5, distinguish outcome distribution from whether those outcomes satisfy the target state. Several classical outcomes may produce a lower CCI without improving post-quantum coverage.

Layer 1's saturated baseline is not, by itself, a new finding every cycle. Its applicable asset count and service distribution describe the remaining classical public-key estate. A layer-1 consequence-threshold breach, an unapproved horizon or a missing migration route can nevertheless be a material finding. The baseline convention must not suppress those consequences.

Layer	High reading	Remediation lever
1	Nothing distinguishing: every institution reads 10,000 until quantum-vulnerable material	Migration only; procurement cannot move it

Layer	High reading	Remediation lever
	retires	
2	One code defect reaches most of the service	Deliberate library diversity; lineage disclosure in procurement
3	Root compromise or deprecation halts the service; the headline may be diversified while the flagged view is not	Dual-root architecture; separate operated from trusted anchors
4	Firmware defect or module recall reaches most key material	OEM diversity at firmware-family level; separate master from bulk keys
5	Downgrade or negotiation failure exposes most paths	Counterparty coordination; remove weak options where better is supported
6	A defective entropy source or generation design is silent, total and retroactive	Design diversity across generation points; evidence chaining to certification records

Layer 1 is the baseline, not a differentiating finding. Layer 6 is not visible in a vendor-layer assessment.

D.2 MATERIALITY

Materiality is assessed from the service effect and the decision required. Give priority to a demonstrated tolerance breach, an upper-bound scenario that the institution cannot currently exclude or accommodate, a critical dependency with no equivalent, and the absence of operational facts necessary to determine whether a severe consequence is tolerable.

A material divergence between the headline and flagged-only views warrants review of which functions the qualifying assets support. A small flagged population naturally produces a high CCI; the asset count and consequence are necessary context. A difference of more than one descriptive band is a useful review trigger, not an automatic proof of unacceptable risk.

Consider remediability explicitly. Where an independent alternative is available, the decision may concern design or procurement. Where it is not, the appropriate response may be additional recovery capability, stronger verification, containment, supplier evidence, coordinated migration or documented risk acceptance. Lack of an available alternative does not reduce the measured concentration.

Low discovery coverage or indirect evidence limits the strength of a conclusion. It does not justify ignoring a directly evidenced critical exposure until every inventory gap is closed. Separate actions to improve evidence from actions supported by facts already known, and distinguish both from speculative changes based only on an uncertain score.

Materiality is read in order: 1 tolerance result – Fail is material regardless of band, Pass at 5,200 usually is not, and Not assessable is material where the missing fact is the threshold or recovery estimate; 2 flagged-only divergence of more than one band; 3 remediability – a lever creates a finding with an owner, while no lever creates a risk-acceptance decision; 4 coverage confidence – a reading over 40% coverage at E5 is a hypothesis. Do not escalate on band position alone.

D.3 CHALLENGE QUESTIONS

For layer 1, verify scope, wrapping dependencies, parameter classification and treatment of hybrids. Ask whether missing classes have been excluded by rule or simply not observed. For layer 2, ask what identifies the executing family, which code remains shared, what establishes divergence, and which unresolved populations could share an unidentified ancestor.

For layer 3, inspect the accepted paths and the verifier's policy, not only certificate chains in a repository. Ask whether an alternative path preserves availability, whether a compromised chain remains acceptable, and whether a policy authority actually has the powers attributed to it. For layer 4, challenge firmware-family grouping, manufacturer provenance and the handling of managed-service boundaries.

For layer 5, identify the observation source and window, the affected configuration population, and any downgraded or unobserved tail. For layer 6, trace generation rather than current storage, distinguish generic design class from shared implementation identity, and review whether an interval is named-node reach or an bounded design.

Across all layers, challenge denominator construction, duplicate removal, service-specific criticality, the basis for exclusion from bounds, recovery assumptions, and any claim that a changed score represents remediation. Review the full assessment supplement where a scalar CBOM property cannot express the required relationship.

Layer-specific challenge questions are mandatory. Layer 1: unobserved asset classes, wrapping keys, and removal of post-quantum material from the denominator. Layer 2: evidence for each lineage, evidenced versus asserted divergence, `UNDISCLOSED` count,

response rate and reach bounds. Layer 3: accepted anchors, cross-signing, alternative paths, operated versus merely trusted anchors. Layer 4: firmware family versus product name, OEM origin, and cloud-module disclosure. Layer 5: outcome versus capability, observation method and sampled paths. Layer 6: design versus certification profile, bounded design spanning conformant points, and sample basis. Cross-cutting: coverage method, tier distribution, finely balanced determinations and what would reverse them.

D.4 REVIEWER ERROR CATALOGUE

Error	Why it matters
Merging distinct executing families merely because they share an ancestor	Confuses the primary-unit CCI with overlapping reach
Counting both hybrid components as full votes for one asset	Breaks the exclusive share calculation and biases the denominator
Calling the largest point's share evidenced reach of an undisclosed design	Converts a mathematical bound into an unsupported provenance claim
Replacing an unknown module with its cloud provider's name	Invents the sublayer identity the assessment is intended to establish
Treating generic certification as proof of independent origin	Misstates the validation's scope and can conceal shared dependencies
Substituting a capability list or roadmap for observed target state	Overstates operational coverage
Treating marginal population percentages as known joint coverage	Can substantially overstate end-to-end readiness
Reporting the largest upstream node while omitting a larger standalone unit	Understates the largest failure domain
Reporting missing or empty populations as CCI zero	Creates artificial diversification
Selecting evidence by code number without checking its claim	Can prefer a stale inventory assertion to relevant primary evidence
Converting a CCI band directly into downtime, loss or a compliance verdict	Omits the operational determination and exceeds the metric's scope

Error	Why it matters
Publishing a bare institution mean or delta	Removes information needed to interpret both the level and its movement

D.5 FEASIBLE DIVERSIFICATION

Supply constraints influence the response to concentration, not the calculation. Where architecture or procurement proposes diversification, record the available independent implementation, firmware or generation options, their provenance, interoperability, validation requirements, operating cost and replacement lead time. An unsupported claim that the market contains only a small number of independent lineages is not an adequate substitute for this assessment.

For sector analysis, publish the scope and evidence behind any census of shared dependencies. Institutional CCI figures cannot establish the market-wide count of independent implementations, and vendor revenue rankings do not establish cryptographic failure-domain shares. Comparable, safely aggregated dependency identities are needed before sector conclusions can be drawn.

At layers 2, 4 and 6 the achievable score is bounded by supply: firmware families and key-generation designs are fewer than the brands selling them. Report a reading above a band boundary with market context – independent dependencies available to buy and independent dependencies held. An institution at 6,000 may hold the most diversified position the market supports, while one at 2,400 may hold the only two independent lineages that exist. Appetite is set against this frontier. The count of independent lineages in the supplier market and the disclosure response rate are themselves concentration figures a supervisor should publish. The proposition that independent post-quantum lineages available to any buyer are "in the single digits" is what the frontier census tests; methodology is at ccframework.org/census, with snapshot one due at v1.0 final. Vendor revenue rankings do not establish cryptographic failure-domain shares.

PART E – REMEDIATION

E.1 SELECTING AN EFFECTIVE RESPONSE

Remediation should identify which dependency, coverage gap or consequence it changes. The target is not necessarily a lower CCI in every cell. An institution may improve end-to-end protection, recovery capability or evidence without changing the distribution of primary units. It may also reduce a primary-unit CCI while leaving a common ancestor's reach unchanged.

Algorithm exposure. Replace or protect vulnerable functions in line with the approved target and horizon. Report the remaining baseline population and coverage rather than expecting layer 1's CCI to decline gradually. A sound accepted hybrid can improve protection while the classical component remains in the baseline.

Implementation lineage. Select and verify distinct implementations where the operational benefit justifies the additional testing and maintenance. Require evidence of relevant code paths, build configuration and update governance. Supplier disclosure improves the resolution of reach; it does not by itself change the deployed code. A change of executing family that retains the same affected ancestor may change the CCI without resolving the common-mode scenario.

Trust roots. Assess independent trust paths, internal/external authority separation, certificate replacement and distribution capability. Dual-root arrangements must be designed for the intended failure mode: availability after root withdrawal and resistance to forged acceptance are different objectives. Adding a trusted root without a corresponding policy may increase rather than reduce the acceptance surface.

Custody. Consider genuinely distinct firmware and hardware domains, portable key-management arrangements where permissible, and segregation of high-consequence key material. Additional platforms introduce cost, operational complexity and new interfaces. The assessment should compare that cost and transition risk with the recovery or containment benefit, rather than mandate diversity for its own sake.

Protocol outcomes. Remove unnecessary below-target options where both parties can operate safely without them. Coordinate the remaining changes with counterparties and verify the outcome after deployment. Configuration changes, certificates, policy and operational onboarding may be as important as algorithm support.

Key generation. Establish generation provenance, replace affected or inadequately evidenced designs where justified, and plan replacement of existing key material where

the scenario requires it. A new generator does not automatically remediate keys generated earlier. Diversification must be assessed at the relevant design and generation-service levels, not by device badge alone.

Hybrid composition is defence in depth under specified assumptions, not a general declaration of independent implementations. A shared component does not automatically invalidate a mathematically sound hybrid; it creates an implementation failure domain that must be tested. Conversely, two independently named algorithm implementations do not establish a secure combiner or deployment.

Layer 1 remediation is migration to post-quantum algorithms and nothing else. No procurement action moves this layer; report progress as estate profile, not as falling concentration. Layer 5 remediation is often free: paths negotiate down because no one verified the negotiated outcome. Layer 6 moves slowly because the estate turns over slowly.

Hybrid reduces consequence, not concentration. The classical component keeps the asset at layer 1. In an AND-compromise construction both components must fail and credit is available through C.6's soundness test. In an OR-compromise construction either component or the combiner defeats the exchange; it earns no credit at any layer and adds a lineage dependency without protection. A hybrid whose components share a lineage has no implementation independence.

E.2 CHANGES THAT MAY NOT ADDRESS THE FINDING

Proposed change	Required challenge
Add a supplier using the same upstream code	Does the relevant ancestor's reach change, or only the executing-family and commercial distribution?
Buy a differently branded HSM	Is the firmware or hardware failure domain actually different?
Add another intermediate CA under the same root	Does a tested failure retain an independently acceptable path?
Install a PQC-capable product	Do the service's paths and counterparties actually select and enforce the target?
Obtain another validation certificate	Does it add relevant provenance or operating evidence, rather than only another certificate reference?

Proposed change	Required challenge
Rotate keys without correcting the generator	Can the same defect recur in the new material?
Improve discovery and declare a lower result remediation	Was there a beneficial estate change, or only changed scope and attribution?
Adding a supplier without checking ancestry	Does the vendor register change while no cell in the heat map changes?

A change that leaves the measured concentration unchanged may still improve recovery, assurance or containment. Such benefits should be reported under the relevant scenario, not dismissed because the CCI is unchanged. Equally, a changed CCI is not sufficient evidence that the original common-mode risk has been addressed.

A proposed remediation that changes the vendor register and no cell in the heat map is a false move.

E.3 SEQUENCING AND ACCEPTANCE

Prioritise intolerable service consequences, critical assets without an equivalent, and changes with long external lead times. Begin engagement with blocking counterparties early, but do not assume every service path is held by every blocker. The dependency map identifies which populations each action can release.

Sequence engineering changes to preserve service continuity and permit verification. Diversifying several cryptographic layers simultaneously may introduce more transition risk than a staged approach. Record testing, rollback restrictions, key/certificate migration dependencies and the basis for accepting residual shared exposure during the transition.

Each material finding must have an action owner, an intended effect, evidence of completion and a review date, or a documented acceptance by an authorised forum. Acceptance identifies the scenario and the limitation accepted; a generic acceptance of "the score" is insufficient.

Three sequencing rules apply. Fix what fails tolerance first, regardless of band. Prefer flagged-only divergence: a layer at 2,400 headline and 7,200 flagged-only is more urgent than one at 6,000 flat, because the concentration is in material whose compromise is irreversible. Treat blocking counterparties as a long-lead item started in the same cycle as the assessment.

PART F – ASSESSMENT MATURITY

Maturity describes how the assessment is operated, not how much technical data exists or how low its scores are. The data levels CCF-1/2/3 and coverage markers EC-0/1/2 are separate attributes.

Level	Discovery and evidence	Resolution and challenge	Governance
M1 – Initial	Partial or tool-scoped sources; completeness may be unstated	Preliminary assignments; limited recorded judgements	One-off exercise or exploratory baseline
M2 – Repeatable	Repeatable collection and a stated coverage method	Documented units and determinations; gaps retained	Named responsibilities and recurring assessment
M3 – Governed	Coverage challenged by second line; material claims graded with exceptions retained	Material determinations challenged; relevant bounds and applicability assessed	Approved use, appetite/acceptance process and model-classification determination
M4 – Assured	Maintained supplier-evidence programme and quality monitoring	Independent review of computation and significant assumptions; technical confirmation of material resolutions	Results inform procurement, resilience and remediation decisions
M5 – Embedded	Maintained inventory and monitored change between	Reusable precedents, event-driven updates and tested reporting	Defined breach actions and sustained use across

Level	Discovery and evidence	Resolution and challenge	Governance
	cycles	controls	institutional programmes

Assign the highest level whose relevant conditions are met across the declared assessment scope, and record material exceptions. A higher data level does not waive the M3 reporting gate. An assessment may be well governed despite suppliers withholding advanced evidence; its limitations and response programme must be explicit.

Improved maturity may change the figures in either direction as discovery and attribution improve. Do not treat an increased CCI as evidence that the assessment process has deteriorated. Evaluate maturity against the process criteria and explain metric movement under C.8.

An institution moving from M1 to M3 will usually see its index rise because better discovery surfaces dependencies that were invisible. Reporting maturity beside the index stops that rise reading as deterioration and removes the incentive to look less hard. Maturity and conformance are different properties: CCF-3 data under an M1 assessment is worth less than CCF-1 data under M3.

PART G – PROGRAMME INTEGRATION

G.1 RISK TAXONOMY AND RCSA

Place cryptographic concentration within the institution's existing technology, cyber, third-party or operational-risk taxonomy with explicit accountability. A separate risk entry is appropriate where the current taxonomy would otherwise conceal the common-mode exposure; creating a duplicative risk category is not required merely to adopt CCF.

An illustrative definition is: "The risk that shared cryptographic dependencies cause correlated impairment of service confidentiality, integrity or availability across apparently separate implementations or suppliers." The RCSA should identify the service effects, preventive and recovery controls, evidence gaps and action owners. CCI and reach are structural indicators within that record, not substitutes for the control assessment.

G.2 RISK APPETITE

Set appetite with reference to service tolerances, criticality, feasible alternatives and institutional risk capacity. The current score distribution can inform implementation priorities, but should not be used simply to set thresholds that no current service breaches. Transitional exceptions can be explicit where an acceptable end state requires time or external coordination.

An illustrative statement is: "Services exceeding [CCI or reach threshold] at layers 2 to 6 require a documented scenario assessment and approval by [forum]. Required target-state coverage for [service/function] must reach [target] by [date]. Material tolerance failures and unbounded critical dependencies must have a named response owner and review date." The institution must specify the intended outputs and their qualifiers; a generic numerical threshold should not silently apply to both CCI and reach.

Risk appetite applies at any of layers 2 to 6. Set the threshold from the institution's own distribution, because an appetite every service already breaches produces no behaviour. Keep CCI and reach thresholds distinct.

G.3 SCENARIO ANALYSIS AND CAPITAL PROCESSES

Use the dependency map to select common-mode scenarios affecting several services at once. State which functions are impaired, which counterparties react, what recovery capabilities remain and which controls preserve confidentiality or integrity. This supports severe-but-plausible scenario design and operational planning.

CCF does not estimate scenario likelihood, monetary loss or regulatory capital. Any use in ICAAP or a separate financial model requires its own transmission assumptions, governance and validation. A common dependency identified by CCF is an input to that analysis, not a quantified capital requirement.

Parameterise the common-mode scenario as failure of the largest dependency at the institution's highest-reading layer, across every service that depends on it simultaneously. Standardised operational-risk capital is calibrated to idiosyncratic loss history, so a forward-looking common-mode event is under-weighted.

G.4 THIRD-PARTY RISK AND PROCUREMENT

Add relevant implementation, firmware, manufacturing and generation provenance to supplier due diligence and concentration assessment. Link these records to contractual services and the business services they support. Keep commercial substitutability and cryptographic commonality as related but distinct questions.

The Supplier Lineage Disclosure Request provides a consistent evidence request and response classification. Use contractual information rights, independent attestation or controlled disclosure where appropriate. Track material changes in provenance during the contract term. CCF supports, but does not replace, the institution's obligations under applicable third-party-risk rules.

DORA Article 28(4)(c) addresses whether an ICT contractual arrangement reinforces concentration risk, while Article 29 addresses substitutability and concentration involving ICT third-party providers. Both operate at the provider and contractual-arrangement level and do not require resolution of cryptographic dependencies beneath them. Extend the assessment with lineage, firmware family and manufacturer origin per critical provider; add lineage disclosure to due diligence and record the response.

G.5 OPERATIONAL RESILIENCE

Add cryptographic dependencies to the existing service map and test common-mode failures against approved tolerances and consequence thresholds. Recovery exercises should consider shared operational constraints: limited replacement capacity, coordinated certificate distribution, counterparties unable to accept a new profile and issued devices that cannot be updated remotely.

The failure-impact record connects structural results to these exercises. Use measured exercise results to replace provisional recovery assumptions where possible, and retain the distinction between tested and estimated capability. `NOT MAPPED` and Not assessable findings are inputs to improving the resilience process, not merely exceptions to a cryptographic report.

G.6 MODEL RISK AND INDEPENDENT REVIEW

The institution must determine whether its CCF implementation falls within its applicable model definition and document that conclusion. The determination depends on the use, degree of judgement and governing policy; this framework does not provide a blanket regulatory classification.

Regardless of classification, document the calculation, input quality, assumptions, departures, limitations and review arrangements using the Model Documentation Pack. Independent review should cover judgement and mapping as well as arithmetic. A simple formula does not make incorrect population construction or unsupported reach assumptions immaterial.

The Concentration Sensitivity Model Concept Note concerns an optional, separate modelling exercise. It is not required for a CCF assessment and its hypothetical network results must not be represented as measured CCF outputs.

The Model Documentation Pack pre-fills these assumptions: asset-count weighting and its declared bias; maximum across layers rather than average; band derivation from largest share; undisclosed dependencies as reach bounds; minimum across paths for coverage. It pre-fills these limitations: the figure moves with discovery quality; flagged-only indices run high on small samples; layers 2 and 6 typically rest on E4 and E5 evidence; no output is a loss estimate.

PART H – SUGGESTED FIRST-CYCLE PLAN

A first cycle should establish a defensible result for a manageable service scope and a route to improve the evidence. Ninety days is a planning example, not a delivery guarantee or a requirement to defer action until a particular data level is reached.

During the first month, approve the charter, select services, confirm ownership and boundaries, and issue the input and supplier requests. Establish the discovery denominator independently of the tooling. Record the target-state definitions, tolerances and external dependencies early enough to avoid building a technically complete inventory that cannot support the intended decision.

During the second month, resolve counted identities, applicable layers and primary units. Begin with the layers supportable from existing observations and inventory evidence. Develop custody, ancestry and generation relationships where evidence permits. Record judgements as they are made, and distinguish a complete calculation over a scoped population from a claim that the entire estate has been discovered.

During the third month, verify the computations, complete the material failure-impact determinations and report the service-level findings. Compute EC only for a supported path population. Apply the M3 gate to any institution-level figure. Agree remediation, evidence-gathering actions and the next review, and integrate supplier provenance questions into procurement so that subsequent cycles do not restart from the same gaps.

Advanced ancestry and generation evidence may remain incomplete at the end of the first cycle. That should be a documented limitation with an owner and a response programme, not a reason to conceal the other results or to treat missing provenance as independence.

PART I – INSTRUMENTS

Referenced by Part B, published separately at ccframework.org.

Instrument	Used in	Purpose
Data request specification ✓	Phase 2	Fields required per conformance level, per service, with the coverage claim format. *First published at v0.6; current edition on the family baseline.*
Supplier lineage disclosure request ✓	Phase 2	Standard questions on implementation ancestry, firmware family, manufacturer origin and entropy design, with response grading. *First published at v0.6; current edition on the family baseline.*
Discovery coverage attestation (the coverage predicate)	External disclosure	Signed in-toto statement attesting how discovery was performed – reference sets by digest, attempted scope, scanner inventory, gap register – so coverage is demonstrable without disclosing the inventory. *Reserved at RC: predicate type URI* `ccframework.org/predicates/coverage/v1`*; specification in v1.1, Q1 2027.*
Determination log template ✓	Phase 3	Fact pattern, treatment, reasoning, evidence tier, challenge status. *First published at v1.0-RC.*
Assessment workbook	Phases 4 to 5	Computation, with reach, its bounds, and the flagged-only views built in. *Publishes September 2026, per the public roadmap.*
Heat map template	Phase 7	Primary reporting format. *Publishes September 2026, per the public roadmap.*
Board report template ✓	Phase 7	Institution-level figure with mandatory companions and limitations. *First published at v1.0-RC.*

Instrument	Used in	Purpose
Model documentation pack ✓	G.6	Assumptions, limitations, validation plan. *First published at v1.0-RC.*

The reference implementation `compute_ccf.py`, the ten canonical test vectors (markdown and JSON) and the end-to-end reference run publish under Apache-2.0 on the Applied Quantum GitHub, linked from ccframework.org; an implementation that reproduces the vectors is conformant.

PART J – GOVERNANCE, LICENCE AND PROVENANCE

Licence. Creative Commons Attribution 4.0, matching the Applied Quantum PQC Migration Framework. Free to use, adapt and share, including commercially, with attribution to Steve Vaile and Marin Ivezic, Applied Quantum.

Canonical source. Published at ccframework.org. Cite as: *The Applied Quantum Cryptographic Concentration Framework, Universal v1.0-RC, Steve Vaile and Marin Ivezic / Applied Quantum, ccframework.org.*

Versioning. Sector extensions align on this baseline rather than versioning independently, which prevents an extension being run against a core it was not written for. The CBOM Profile versions independently, because it is a shared dependency of two frameworks. The Profile owns property semantics; the assessment supplement owns contextual determinations not fully expressible in those properties.

Corrections. A published correction route, with corrections issued additively rather than by silent edit.

Disclosure. Applied Quantum sells post-quantum migration advisory services, and this framework concludes that cryptographic concentration is a systemic exposure requiring an accountable owner. That interest is stated on the first page of every derived publication. The framework is published openly, the metrics are computable without engaging Applied Quantum, and the assessment can be performed by any institution using its own inventory and its own staff.

Package map and precedence. The family comprises: this Universal Framework – normative method, metrics and reporting; the sector extensions – normative only within their enumerations, never altering a core rule; the CBOM Conformance Statement – the normative data-to-computation contract; the Applied Quantum CBOM Profile – normative data model, controlling property names and value types, versioned independently; the Technical Companion and the whitepapers – informative; the instruments – operational templates, joined at RC by the Discovery Coverage Attestation reservation; the Concentration Sensitivity Model concept note – informative, optional, classified separately under G.6. Where documents conflict, that order of authority resolves it within each document's domain.

Provenance. Adjacent-work survey, first seeded entry – Kemp, A. C., R. Padbury and C. Reeves (KPMG LLP), *The Role of the Second Line of Defense in Post-Quantum Cryptography Migration and Cryptographic Auditing*, SSRN, 2026. Scope: second-line responsibilities for PQC migration and integration into SOX and SOC programmes. Vendor treatment is limited to SLA verification and PQC readiness assessment. No below-vendor dependency, implementation-lineage, key-generation-design or concentration concept appears. The full adjacent-work survey – prior-art positions, method, and the update route – publishes as a standing page at ccframework.org/provenance; entries are added there rather than here.

Original concepts introduced here: the six-layer below-vendor decomposition; the index-and-reach separation, with non-disclosure carried as a reach bound and its spread; the flagged-only index; Effective Coverage as a path-complete migration measure; counterparty classification by remediation route; the conformance-and-maturity separation.

ANNEX 1 – VALIDATION AGAINST THE PUBLIC RECORD: ROCA (2017)

A pilot validates that the framework can be executed. It cannot validate that the layers read the right thing, because a pilot estate has no disclosed failure to check against. The public record does. This annex re-runs the layer-6 and layer-2 readings over the best-documented below-vendor failure on record and states which parts of the result were computable from evidence available before disclosure. Debian (2008, layer 6) and KyberSlash (2023, layer 2) join at v1.1.

THE EVENT

On 16 October 2017, researchers at Masaryk University's CRoCS laboratory, Enigma Bridge and Ca' Foscari University disclosed CVE-2017-15361: the RSA key-generation routine in Infineon's RSALib produced structured primes, making the resulting keys practically factorable from the public key alone. The disclosure page priced the attack on a 2014-era CPU at roughly 2 CPU-hours for a 512-bit key, 97 CPU-days ($40–$80) for 1024 bits, and 140.8 CPU-years ($20,000–$40,000) for 2048 bits. The library shipped inside smartcards, security tokens and trusted platform modules across product lines from many vendors – Microsoft, Google, HP, Lenovo and Fujitsu all shipped mitigations at disclosure – in chips manufactured from 2012, holding FIPS 140-2 and Common Criteria EAL5+ certifications. The vendor had been notified in February 2017; the paper followed at ACM CCS on 2 November.

On 2 November 2017 the Estonian government endorsed suspending the certificates of 760,000 ID cards. The suspension took effect at midnight on 3 November, with renewal through Police and Border Guard service points and online.

THE RETRODICTION

Put an assessment in 2016 at an institution running a TPM-attested device fleet and an issued-token estate.

Observed diversity is high. Several laptop manufacturers, several token brands and several card products mean a vendor-substitutability test passes.

Layer 6 resolution collapses it. The generation point is the embedded chip. Vendor documentation and certification records naming the Infineon chip families and their on-chip RSA generation resolve a brand-diverse estate toward one generation design. Where the record evidences the design, resolution is direct. Where it evidences only profile conformance, the design node is bounded. Its reach spans most or all of the issued estate.

The same event reads at layer 2. The library is also an executing implementation family. One failure appearing in two layers is expected behaviour, not double counting: the layers ask different questions of the same estate.

The determination fails in advance. C.5.2's layer-6 scenario requires re-keying the affected estate and accounts for retroactive exposure of keys already generated. For an issued physical estate, re-keying is a months-long campaign against a tolerance measured in hours or days, and the retroactive term engages the consequence threshold regardless of the clock. Estonia carried out that determination during the event.

WHAT THIS VALIDATES

No pre-2017 evidence could have shown the design was broken; the framework produces no cryptanalytic judgement. The validation claim is only that the framework's outputs – one design behind many brands, reach at or near the whole estate, a re-key time longer than any tolerance, certifications attesting correctness while saying nothing about provenance – were computable from pre-disclosure inputs. The event was chosen because its record is rich, which creates a selection effect.

Validation continues on three tracks: retrodictions, public-population computations through the census, and the reproducibility kit. The pilot validates executability, not the mathematics.

SOURCES, TIERED

Claim	Source	Record class
CVE, mechanism,	CRoCS disclosure page and	Peer-reviewed paper and

Claim	Source	Record class
factorisation costs, affected device classes, certifications, disclosure timeline	CCS'17 paper page, public disclosure	researchers' primary disclosure
760,000 ID-card certificates suspended 3 November 2017 at 24.00; renewal route	Estonian government announcement, 2 November 2017	Government primary record

ANNEX 2 – CANONICAL ASSET IDENTITY AND DE-DUPLICATION

A2.1 COUNTING PRINCIPLE

This annex is normative from this edition: an assessment states conformance to it or states its departures, and C.8's common estate is defined over identities matched under it. An asset is identified by the failure it shares: copies of one key share one compromise, and one certificate is one revocation however many servers present it.

Count the persistent asset represented by the evidence, not each discovery-tool observation or each descriptive component attached to it. State the asset classes included in the assessment and their applicability by layer. The same cryptographic relationship can affect several assets, but a repeated record of one asset must not create additional weight.

Asset class	Identity basis	Required treatment
Persistent key or key pair	Cryptographic fingerprint where appropriate, or a scoped stable platform identifier including version	Collapse replicas, backups and duplicate public/private observations of the same key-pair lifecycle; retain genuinely distinct keys
Certificate	Certificate fingerprint, or a verified issuer identity plus normalised serial number	Collapse repeated installations of the same certificate; issuer display name alone is insufficient where different issuers share a name

Asset class	Identity basis	Required treatment
Cryptographic endpoint/configuration	Stable endpoint or configuration identity, service role and cryptographic function	Record versions and outcomes as state; collapse equivalent operational replicas under a documented grouping rule and record the replica count
Issued-estate assets	Counted instances of a specified key, certificate or endpoint class, with evidenced cardinality	A cohort represents its actual distinct instances; a shared key or certificate copied into many devices is not automatically many distinct keys or certificates

A certificate and its subject key can be separate lifecycle assets when both classes are in scope. Their dependency relationship must be retained, and the mixed denominator must be declared; counting both is not a statement of independent business value. Do not additionally count their algorithm description or path-state record as another protected asset.

Ephemeral per-session secrets are excluded from the persistent denominator. The endpoint/configuration and persistent material responsible for their security remain in scope. Cohort weighting requires a known count and a defensible homogeneous state; an issued population with no cardinality is a coverage gap, not a single asset.

A2.2 MERGING SOURCES

Where two discovery sources report material with one failure identity, record one asset. The higher evidence tier wins. An equal-tier conflict goes to the determination log. The losing value is retained as a recorded discrepancy rather than deleted. A tool that consistently loses merges is itself a finding about the tool.

Platform identifiers are scoped to the issuing system and key version. Do not export secret material to resolve identity.

A2.3 IDENTITY ACROSS ASSESSMENTS

Match assets across cycles using cryptographic identity first, then a demonstrably stable scoped platform identity, and only then a location-and-role match marked `identity-inferred`. Endpoint/configuration identities should remain stable across an

ordinary policy change; record the changed outcome or implementation as state, rather than unintentionally converting every configuration update into asset retirement.

Re-keying creates new key material and is recorded as retirement plus deployment where the old material leaves scope. Matching on location alone would record a re-key as continuity and understate real movement, which C.8 exists to prevent. Renewed or reissued certificates are new certificate assets. Old material that remains relevant for verification, archived confidentiality or another in-scope function is not automatically retired because a newer key or certificate exists.

The common-estate share in C.8 uses these identities and records the matching basis. Material changes in identity or grouping policy require a restatement or an explicit comparability limitation.

A2.4 DIVISION OF RESPONSIBILITY WITH THE PROFILE

The Profile references native identifiers and defines the added properties used to connect the records. This annex defines how CCF counts them. `bom-ref` is a document reference, not by itself proof of canonical identity across tools or cycles. The assessment supplement retains canonical aliases, layer applicability, cohort cardinality and contextual relationships that the current Profile cannot fully encode.

The factual-inventory and assessment-overlay split stays open by design, matching the Profile's identical question, and is decided jointly on pilot input at v1.0 final.

Future identity conventions in the Profile must be checked against this annex before adoption. No new Profile property is implied merely because the assessment requires a counting fact.

RESOLVED AT V1.0-RC

Resolved in v0.7 following review: the disclosure response rate is mandatory wherever requests were issued; the institution-level figure is withheld below M3; the determination log and assessment workbook are the blocking instruments.

Four questions carried from v0.9 review, closed or scheduled August 2026.

1. First cycles run five to eight services, recorded in Part H, and five to twelve remains the steady-state range.
2. Canonical asset-identity and de-duplication rules are normative at v1.0. Annex 2 carries them, with its counting rules coordinated with the CBOM Profile.

3. Consequence thresholds stay the institution's determination under C.5.3. Non-normative threshold templates publish in the Payments and Financial Services extensions.

4. The factual-inventory and assessment-overlay split stays open by design, matching the CBOM Profile's identical question. One answer, decided jointly on pilot input, is published in both documents at v1.0 final.