

AUGUST 2026
APPLIED QUANTUM

THE APPLIED QUANTUM CCF INSTRUMENT MODEL DOCUMENTATION PACK



Assumptions, limitations and validation plan for the CCF computation under G.6

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.

COPYRIGHT AND LICENSE

© 2026 Steve Vaile and Marin Ivezić / Applied Quantum. This work is licensed under Creative Commons Attribution 4.0 International (CC BY 4.0). You are free to share and adapt this material for any purpose, including commercial use, provided you give appropriate credit to Steve Vaile, Marin Ivezić and Applied Quantum, link to the licence, and indicate any changes made.

Licence details: <https://creativecommons.org/licenses/by/4.0/>

The reference implementation and canonical test vectors publish separately under Apache-2.0 on the Applied Quantum GitHub. Neither licence reads onto the other.

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	CCF instrument – model documentation pack
Parent document	The Cryptographic Concentration Framework – Universal – v1.0-RC (August 2026)
Intended audience	Model risk management and validation functions reviewing the computation
Assumed knowledge	The Universal Framework's metric definitions and the reference implementation
Scope	Documentation sufficient for model-risk review of the computation and its inputs: units, exclusivity, reach semantics and the canonical test vectors.

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 published.
1.0-RC	8 September 2026	Validation plan extended with computation, data and judgement tests and an acceptance table; the computation tests align to the ten canonical vectors and the rulings of 8 September. Prose revised throughout.

HOW TO USE THIS DOCUMENT

This pack documents the computation for model-risk review. It pairs with the Universal Framework's metric definitions and the canonical test vectors, and the reference implementation reproduces every figure it documents.

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

1. Classification determination.....	5
2. Section A – Description.....	5
3. Section B – Assumptions register.....	6
4. Section C – Limitations and input quality.....	8
5. Section D – Validation plan.....	9
5.1 Computation tests.....	9
5.2 Data and judgement tests	11
5.3 Acceptance and unresolved findings.....	11
6. Monitoring and change control	11
7. Review record.....	12

Used with G.6 of the [CCF Universal Framework](#) v1.0-RC. Whether the CCF computation is a model is the institution's determination under its own regime and definition – regimes such as SR 11-7 or PRA SS1/23 where they apply. Where the answer is yes, this pack is the required documentation. Where it is no, the same pack is good operating discipline.

1. CLASSIFICATION DETERMINATION

Complete this pack for the institution's implementation of the CCF measurement process. It records how the method is implemented and controlled, whether or not the institution classifies that implementation as a model. The classification is determined under the applicable institutional and regulatory definition, not solely by the simplicity of the HHI formula.

The optional concentration sensitivity model is a different use case. If developed or used, it needs a separately identified documentation pack covering its network, exposure and propagation assumptions.

Regime and definition applied: [instrument, section]

Determination: [model / not a model]

Reasoning: [why, against the definition's elements]

Registered: [model inventory ref, where applicable]

Owner and date: [name, date]

The determination is also a determination-log entry, cross-referenced by ID.

Also record the implementation name, intended use, accountable role, approval authority and date, and any control-register reference.

A pending classification does not justify publishing an institution-level indicator as if the M3 governance conditions were complete. Independent review arrangements must be defined even where the final determination is "not a model".

2. SECTION A – DESCRIPTION

The computation measures cryptographic concentration and Effective Coverage per important business service under Universal Parts A and C. It consumes a CBOM conforming to the Applied Quantum CBOM Profile v1.0-RC or later at a stated conformance level, with per-value evidence tiers and a coverage statement per the Data Request. It produces the CCI over exclusive executing units, failure-domain reach with evidenced and bounded states, flagged-only views, Effective Coverage and the C.8 movement decomposition. The heat map and board report carry these outputs.

Document the assessment services, entities, users and decisions supported. Name the exact Universal Framework, Profile, Conformance Statement, sector extensions and calculation implementation. State whether the implementation produces complete or qualified service maxima, reach bounds, bounded design results, criticality views, EC and movement decomposition.

Describe the input pipeline: source systems, snapshot controls, canonical identity, population weights, service mappings, evidence grading, supplier records and the assessment supplement. Identify manually maintained transformations and judgement points. A box labelled "CBOM" is not a sufficient description where the calculations require additional path or cohort relationships.

Document the output controls, including rounding, not-assessed/non-applicable handling, interval propagation, the M3 institution-reporting gate and the mandatory companions. State prohibited uses, including unqualified cross-institution ranking and monetary translation of CCI.

3. SECTION B – ASSUMPTIONS REGISTER

Pre-filled from Universal C.7 and G.6. Record acceptance or a reasoned departure for each row, with an owner and the implementation evidence or limitation.

Assumption	Rationale, from the framework	Acceptance / departure	Owner
Asset-count weighting, with its declared bias	The ceiling of what institutions can produce; transaction-volume weighting would be a specification no institution could compute		
Maximum across layers, not average	The worst layer is the failure mode; a mean would conceal it		
Bands derived from largest share	Defect propagation, not market power		
Undisclosed dependencies carried as reach bounds	Non-disclosure widens a bound and never changes the		

Assumption	Rationale, from the framework	Acceptance / departure	Owner
	index		
Minimum state across paths for coverage	Path semantics preserved inside an asset denominator		
Layer applicability	Each cell uses its applicable deduplicated population	[acceptance/departure; rules and excluded/unobserved classes]	[owner]
Exclusive primary assignment	One vote per asset; composite units where necessary	[acceptance/departure; resolution and component-role records]	[owner]
Shared dependency reach	Non-exclusive asset sets with evidenced and bounded membership	[acceptance/departure; relationship graph and candidate exclusions]	[owner]
Unknown designs	For a profile-only design node, known reach is the largest conformant point's own design, a floor by construction; the bound spans all conformant points without establishing that they share a design	[acceptance/departure; profile evidence and unresolved population]	[owner]
Criticality	Five service-specific tests; unknown does not equal false	[acceptance/departure; owner determinations and unresolved counts]	[owner]
Coverage	All relevant functions/paths; joint population evidence or bounds	[acceptance/departure; topology, states, target definitions and sample basis]	[owner]

Assumption	Rationale, from the framework	Acceptance / departure	Owner
Roll-up	Maximum across layers 2–6; service-exposure-weighted institution mean when permitted	[acceptance/departure; weights, missing inputs and gate]	[owner]
Movement	Defined counterfactual ordering and cause records	[acceptance/departure; identity matching and reconstruction limitations]	[owner]

Persistent asset counts include evidenced cohort cardinality; retain the counting and reconciliation records. Bands are descriptive CCI conventions, not calibrated security thresholds; retain the reporting wording.

Record departures explicitly. A different weighting or a changed primary-unit definition is not an unqualified CCF calculation merely because it is documented here. Provide the normative result as the primary output where possible and label alternative calculations as supplementary. Where a required rule is not implemented, state the effect on the conformance claim and use.

4. SECTION C – LIMITATIONS AND INPUT QUALITY

Discovery quality changes the figure, and the institution decomposes that movement under C.8. Flagged-only indices run high on small samples. At layers 2 and 6 the institution typically grades the evidence E4 and E5, the strongest routinely admissible basis there, not a deficiency. No output is a loss estimate.

Describe the actual limitations of this assessment, not only generic caveats. Include undiscovered or unquantified classes, uncertain identities, unsupported ancestry, mismatched validation environments, missing path states, sparse joint-population evidence and unverified recovery estimates.

Report the achieved CCF level and EC marker by scope, the discovery fractions with their denominator statements and the coverage statement's named accountable individual, criticality completeness, the evidence-tier distribution and evidence by material claim class, and the supplier disclosure register where used. Absence is never a favourable value. Explain how each material limitation affects a specific output or determination. E-codes must not be multiplied into shares or used as automatic confidence weights.

A CCI describes the partition of observed primary units and is sensitive to identity and grouping rules. Reach may be a conservative possibility range rather than an empirically estimated distribution. The largest share of one generation point is not evidence of an unidentified common design. A high flagged-only CCI on a small population requires its denominator and operational consequence to be reported.

5. SECTION D – VALIDATION PLAN

Validation covers arithmetic, data transformation and judgement. The reviewer must be sufficiently independent of the implementation and its material assumptions; organisational placement follows the institution's arrangements rather than automatically being assigned to internal audit.

1. **Computation verification.** Recompute the canonical test vectors independently and match the published fixtures; the reference implementation and vectors exist for exactly this. Then recompute a sample of the institution's own cells.
2. **Determination sampling.** Sample entries for every determination type listed in Section 2 of the log template. Each entry must be defensible from its recorded reasoning alone.
3. **Coverage re-performance.** Re-perform the coverage statement's method challenge for the largest service.
4. **Movement audit.** From the second cycle onward, re-derive the C.8 decomposition from the recorded causes for one service.
5. **Review triggers.** Review any departure from Universal Annex 2, any tolerance change and any cycle where the measurement effect exceeds real movement. Review annually regardless.

5.1 Computation tests

Test	Expected result or invariant
One primary unit carrying all assets	CCI 10,000; largest share 100%
Two equal units	CCI 5,000; largest share 50%; no inference of a 70% largest share
Four equal units	CCI 2,500, in the Concentrated band
Exact boundary 8,100	Single-source-equivalent; not the lower band
Empty applicable population	Not applicable, not numerical CCI zero

Test	Expected result or invariant
Counts 2,400 / 900 / 900	CCI 4,183.673469..., presented as 4,184
Named reach 2,400 known plus 900 possible out of 4,200	57.142857...% known; 78.571429...% upper; no change to correctly resolved primary counts
Profile-only design node over 2,650 / 1,050 / 500	Known reach 63.095238...% (the largest point's own design, a floor); bound 100%
Two marginal target proportions 80% and 70% on the same cohort	Joint coverage bounded 50% to 70% without overlap evidence
Same asset copied by two tools	One counted asset after supported identity matching
Hybrid with two simultaneous implementations	One primary vote under the composite rule; both component relationships retained in reach, with both implementation lineages reaching the asset in full
Prior 6/4, current 8/4 with one new and one discovered A asset	Prior 5,200; counterfactual 5,371.900826...; current 5,555.555556...; components sum to delta before rounding
Diversity ladder's largest failure domain	The larger of the largest primary-unit share and the largest reach upper bound; name the endpoint and whether it is a standalone unit, a named upstream node or a bounded design
Unresolved boundary with a known population	Count as one primary unit at its known population; retain uncertainty in reach, with no primary-identity interval test

The reference implementation `compute_ccf.py`, the ten canonical test vectors and the reference run publish under Apache-2.0 on the Applied Quantum GitHub, and published fixtures are pinned by revision. Test the institution's implementation and retain results, including floating-point/rounding tolerances and negative cases. Investigate any conflict between the published fixtures and the applicable method rather than assume equivalence.

5.2 Data and judgement tests

Reconcile a sample of computed cells to source records and repeat de-duplication for material asset classes. Verify that algorithm/library/path descriptions do not add unwanted weight, that key generation is not inferred from current storage, and that relationships remain resolvable after format conversion.

Sample material determinations across the categories in the log. Include primary-family versus ancestor treatment, composite units, cloud boundaries, trust alternatives, criticality, coverage joint states, upper-bound failures and unknown consequence thresholds. A judgement must be supported by source facts as well as by a well-written rationale. Check any named supplier response states against the Profile enumeration.

Re-perform the discovery-coverage claim for material services and challenge the independently enumerated denominator. Verify that a scoped percentage is not presented as whole-estate coverage and that EC-0 is not displayed as 0% target-state coverage.

Test the reporting gate and mandatory companions. The largest-reach result must include primary standalone domains. Unknown EC values and possible band crossings must not disappear from the institution summary. From the second cycle onward, reproduce the movement counterfactual and the two common-estate shares.

5.3 Acceptance and unresolved findings

Test / scope	Reviewer	Result	Finding / severity	Resolution or acceptance	Retest date
[test]	[reviewer]	[result]	[finding]	[decision]	[date]

State the validation scope and tests not performed. Do not describe a schema check as validation of source truth, or an arithmetic test as validation of the framework's predictive accuracy. Material unresolved findings must limit the approved use or be accepted by the appropriate authority.

6. MONITORING AND CHANGE CONTROL

Review at least annually and after any departure from Universal Annex 2, any tolerance change or material changes in the method, Profile, counting policy, service boundary, target, important supplier provenance or calculation implementation. Review any cycle where the measurement effect exceeds real movement. Active migration may warrant more frequent review. A material measurement effect relative to estate movement is a trigger to examine comparability, not proof that either movement is adverse.

Version-control the calculation, transformations and evidence records. Retain approval, test and rollback records. Monitor unresolved data, stale supplier statements, unquantified asset classes and the completion of agreed actions. Revalidate the affected outputs when a change alters semantics rather than only presentation.

7. REVIEW RECORD

Date	Reviewer	Scope	Findings	Resolution	Independence	Versions	Approved use / restrictions	Next review
[date]	[reviewer]	[scope]	[findings]	[resolution]	[independence]	[versions]	[decision]	[date]

The completed pack is referenced by the assessment report and retained with the determination log. A concise implementation may need a concise pack, but none of the essential questions – what was measured, on what evidence, with which assumptions and under whose approval – should be omitted.
