AETHON CREDIT INTELLIGENCE

ACI Framework v1.0

Risk Methodology

Full Criterion-by-Criterion Edition · May 2026 · Version 1.0
Public — aethoncredit.com/methodology

Document	ACI Risk Methodology
Document Version	1.0 — Full Criterion-by-Criterion Edition
Framework Version	ACI Framework v1.0
Framework Active Since	May 2026
Next Review	May 2027 or material event trigger
Classification	Public — published at aethoncredit.com/methodology
Determinism	Same inputs always produce the same output

QUANTITATIVE RISK ANALYTICS OUTPUT

1 GLOSSARY OF DEFINED TERMS

The following terms are used throughout this document and across the ACI platform with the meanings set out below.

Term	Definition
ACI	Aethon Credit Intelligence. The platform providing deterministic, methodology-driven risk indicators for digital asset credit markets. Eight modules defined. The stablecoin module (M1C) is structured into two sub-modules: M1C-A (CeFi) and M1C-B (DeFi).
ACI Framework v1.0	The currently active CRO-approved version of the deterministic computational framework. Active since May 2026.
ACI Risk Indicator	The integer composite score in [0, 100] computed by the engine. Produced under a defined computational framework and not issued as a credit rating.
ACI Risk Band	One of four classifications: LOW (80–100), MEDIUM (60–79), ELEVATED (40–59), HIGH (0–39).
ALGO_STABLECOIN_CAP	Hard cap limiting the final score for any algorithmic stablecoin to 20.
Append-only	DB trigger discipline blocking UPDATE and DELETE on score_snapshots and related tables.
Bucket Score Map (BSM)	Lookup table mapping each bucket value within a criterion to a score in [0, 100].
Cascade Penalty	Engine rule subtracting five points from the raw score when three or more criteria score below 40.
CeFi	Centralised Finance. Covered under Module 1C-A.
Composite Criterion	A scoring criterion whose score is derived from a weighted sub-formula across two or more sub-components.
Convexity Classifier	Display label — POSITIVE / NEUTRAL / NEGATIVE — describing structural payoff shape.
CRO	Chief Risk Officer. Methodology approval authority.
DeFi	Decentralised Finance. Covered under Module 1C-B.
Depeg	Deviation of a stablecoin's market price from its target peg. Tracked over a 90-day rolling window.
Determinism	Identical inputs produce byte-identical outputs across processes.
Duration Multiplier	Scalar applied to the raw weighted score to reflect compounded risk at longer durations. Module-specific.
Hard Cap	Engine-level override applying a mandatory ceiling, floor, or band reclassification.
HIGH_DEPEG_OVERRIDE	Hard rule forcing HIGH band for any stablecoin with 90-day max peg deviation > 1.5%.

HV30	30-Day Historical Volatility. HV30 > 0.35 caps market-risk criterion at 5.
HV1Y	252-Day Historical Volatility. Primary sub-component (50%) of M1B Market Risk composite.
IC Package	Investment Committee Package — structured risk output for investment committee review, comprising score, evidence, stress scenarios, methodology citation, and audit trail. Paginated PDF deliverable from a locked snapshot.
Linear Interpolation	Scoring method for continuous-value criteria — smooth between bucket boundaries.
Methodology Version	A formally CRO-approved framework state. Every score references its computation version.
Module	One of eight scoring modules under ACI Framework v1.0. The stablecoin module (M1C) comprises two sub-modules: M1C-A (CeFi) and M1C-B (DeFi).
OpenTimestamps	Standard for anchoring SHA-256 hashes on Bitcoin's blockchain.
Proof of Reserves (PoR)	Independent verification of platform-held assets, typically by a named auditor.
Reference Scenario	The default analytical starting point presented when a user opens a module. Populated from the user’s own inputs — ACI does not pre-populate assumptions. The Reference Scenario anchors the IC Package output and serves as the baseline against which stress scenarios are compared. It does not constitute a recommendation or model portfolio. Auto-selected starting allocation.
SDACR	Senior-Debt-Adjusted Coverage Ratio. M1B primary credit signal.
SHA-256	Cryptographic hash function used for snapshot content hashing.
Snapshot	Immutable record of a single computed score, persisted append-only.
Stress Engine	Module-specific deterministic engine projecting outputs against pre-defined adverse scenarios.
TVL	Total Value Locked. Aggregate value of assets deployed within a CeFi platform or DeFi protocol.
VOLATILITY_REGIME_ACTIVE	Hard cap activating when HV30/HV1Y regime multiple > 1.5×.
Worst-Case Protocol	Substitution of CRO-approved conservative scores when criterion confidence < 0.70.

2 SERVICE DESCRIPTION AND INDEPENDENCE

2.1 What ACI Is

ACI provides quantitative analytics and scenario computations designed to support independent analysis of digital asset yield-bearing positions. Outputs are generated through a deterministic computational framework using publicly available data and user-defined inputs. Each output reflects the application of a defined methodology, not the opinion of an individual analyst. ACI outputs are computational analytics artifacts and are not intended as investment recommendations, opinions, or credit ratings within the meaning of any applicable regulatory framework. Primary output types produced by the platform are: Score Snapshot (a point-in-time deterministic risk indicator with SHA-256 content hash); IC Package (a structured investment committee document comprising score, criterion breakdown, stress scenarios, methodology citation, and audit trail); Scenario Package (one or more user-defined scenarios with full parameter transparency); and Verification Output (a publicly accessible cryptographic confirmation of any snapshot by snapshot ID).

ACI provides quantitative analytics outputs to support independent analysis. Users determine how outputs are applied within their own analytical and governance processes.

2.2 Independence

ACI operates independently of the entities it evaluates. The platform does not accept compensation from any provider, issuer, or protocol in connection with the production, modification, or publication of outputs. All indicators are generated exclusively from publicly available data under the published methodology.

ACI maintains policies designed to prevent conflicts of interest, including restrictions on employee holdings in scored entities. Outputs are the result of deterministic computation and are not influenced by issuer interaction.

> **Independence Statement**
>
> ACI holds no financial interest in any entity, instrument, or allocation referenced in any document it produces. Risk indicators are deterministic outputs of an algorithm applied to public evidence — not the opinion of any individual analyst, and not the result of any issuer engagement.
>
> ACI maintains policies designed to prevent conflicts of interest, including restrictions on employee holdings in scored entities.

2.3 How to Use This Document

This document is the complete public reference for ACI Framework v1.0. It supports four distinct use cases:

Reproducibility — any ACI risk indicator can be independently replicated using this methodology and the same publicly available inputs.

Score traceability — every indicator on the platform carries a reference to its framework version and a SHA-256 content hash.

Audit-grade verification — institutional users, external auditors, and regulators can verify any score by snapshot ID via the public score verification endpoint (technical details at aethoncredit.com/methodology/verify).

Append-only versioning — scores produced under a prior framework version remain permanently retrievable. No retroactive revision.

3 EXECUTIVE SUMMARY

3.1 The Structural Gap This Methodology Addresses

Digital asset yield-bearing positions introduce structural risks that are not captured by traditional financial frameworks. These risks include on-chain collateral dynamics, smart contract vulnerabilities, governance exposure, and stablecoin peg behaviour. Existing analytical frameworks were not designed to evaluate these factors in a systematic and reproducible way.

The 2022 collapse of Celsius, BlockFi, and Voyager — followed by FTX later that year — was not primarily a market event. It was a structural failure: institutional capital allocated to digital asset yield-bearing positions without access to the analysis those positions required. No existing framework possessed methodology equipped for on-chain reserve verification, BTC-collateralised lending mechanics, stablecoin peg behaviour, or DeFi governance attack surface. The market substituted reputation, yield, and trust. The outcome was predictable.

ACI Framework v1.0 addresses this gap by defining a quantitative, transparent, and deterministic approach to risk measurement. The objective is not to replace judgment, but to provide a structured analytical foundation for it.

Bottom Line Up Front

ACI provides deterministic, transparent, independently replicable risk indicators across eight modules, including a stablecoin module structured into two sub-modules: Phase 1 (BTC-collateralised lending, BTC treasury preferred shares, and stablecoin yield — the latter structured into CeFi (M1C-A) and DeFi (M1C-B) sub-modules) and Phase 2 (mining and validator infrastructure, market-neutral strategies, digital asset venture, tokenised real-world assets, volatility positions).

Every indicator is computed from public evidence via a published, version-controlled formula. Every snapshot is SHA-256 hashed. No issuer pays for scores. ACI maintains policies preventing financial interest in any scored entity.

3.2 Empirical Calibration

Criteria and thresholds were calibrated against reconstructed historical stress events and public failure cases spanning 2020 through 2025. The dataset covers CeFi platform insolvencies (Celsius, BlockFi, Voyager, FTX, Genesis), DeFi protocol exploits (Euler, Cream, Beanstalk, Mango Markets, Wormhole), stablecoin failures (TerraUSD, Iron Finance, USDN), operational stress events (USDC during the SVB stress in March 2023), Bitcoin mining infrastructure failures and distress events (Core Scientific, Compute North, Argo Blockchain, Iris Energy, Stronghold, Greenidge, Mawson, Hive Digital, Marathon Digital, Bitfarms), and market-neutral strategy failures (Three Arrows Capital, Alameda Research, BlockFi, Babel Finance, Genesis Trading, Mango Markets). Criteria with strong discriminative power received higher weights; criteria that failed the discrimination test were reduced. Historical reconstruction testing across the observed failure dataset identified all

reconstructed failures as ELEVATED or HIGH at T-30. Results remain subject to dataset limitations and ongoing recalibration.

3.3 What Distinguishes ACI

Deterministic, code-enforced — same inputs always produce the same output. Verified by SHA-256 determinism tests covering all currently active scoring modules; a determinism regression is a critical control exception that prevents release.

Append-only persistence — scores are immutable once computed. Database triggers block UPDATE and DELETE at the engine level, including for service_role keys.

Optional Bitcoin anchor — content hashes may be anchored on Bitcoin's blockchain via OpenTimestamps for trust-minimised external verification.

Eight modules defined, including a stablecoin module (M1C) structured into two sub-modules — coverage spans BTC-collateralised lending through tokenised real-world assets and short-volatility positions.

Convexity-aware — each module attaches a structural payoff classifier (POSITIVE / NEUTRAL / NEGATIVE) to surface tail-risk profiles separately from the composite score.

Stress engines — every module ships a deterministic stress projection (Phase 1: parametric curve analysis; Phase 2: multi-scenario stress engines).

4 MARKET CONTEXT AND STRUCTURAL GAP

4.1 Why This Methodology Exists

Traditional credit rating architecture — Fitch, S&P, Moody's — was designed for a different universe. Their frameworks depend on audited financial statements, private issuer disclosures, and decades of default data from traditional finance. That architecture serves sovereign debt and corporate bonds well. It was never designed for BTC-collateralised lending, stablecoin yield protocols, BTC treasury preferred equity, or on-chain real-world assets. The risk factors that govern outcomes in digital asset yield-bearing positions — on-chain reserve verification, collateral liquidation mechanics, stablecoin peg behaviour, DeFi governance attack surface — are structurally absent from traditional agency methodologies.

ACI was built to close that gap. The table below sets out the key structural differences.

4.2 ACI vs Traditional Credit Rating Agencies

Dimension	Fitch / S&P / Moody's	Aethon Credit Intelligence
Asset coverage	Corporate bonds, sovereign debt, structured products	BTC lending, stablecoin yield, BTC treasury preferreds, infrastructure, market-neutral, venture, tokenised RWA, volatility
Primary data source	Audited financials, private issuer disclosures	Public on-chain data, regulatory registries, prospectuses, real-time market APIs
Output scale	Letter grades — analyst consensus	Numeric 0–100 — deterministic formula with published weights
Transparency	Methodology published; issuer inputs confidential	Every weight, formula, sub-criterion published in full; reproducible from disclosure

Conflict of interest	Issuer-pays model	No issuer payment; public evidence only
Update frequency	Annual or event-driven; may lag by months	Market data weekly; structural changes within 30 days of material event
BTC collateral risk	Not assessed	Module 1A primary criterion (Collateral Control, 35%)
Stablecoin peg risk	Not assessed	Continuous monitoring; HIGH_DEPEG_OVERRIDE auto-fires above 1.5% deviation
DeFi protocol risk	Not assessed	Native: Module 1C-B Protocol Security 40%, Governance Risk 30%
Audit / verification	Internal review only	SHA-256 content hash + optional Bitcoin anchor + public verify endpoint

5 PLATFORM ARCHITECTURE

5.1 Empirical Grounding

ACI's calibration is anchored to reconstructed historical stress events and public failure cases from January 2020 through December 2025. The dataset spans four asset classes: CeFi lending and yield platforms (Celsius, BlockFi, Voyager, FTX, Genesis, Babel Finance), DeFi protocol exploits and stablecoin failures (Euler, Cream, Beanstalk, Mango Markets, Wormhole, TerraUSD, Iron Finance), Bitcoin mining infrastructure failures and distress events (Core Scientific, Compute North, Argo Blockchain, Iris Energy, Stronghold, Greenidge, Mawson, Hive Digital, Marathon Digital, Bitfarms), and market-neutral strategy failures (Three Arrows Capital, Alameda Research, Genesis Trading). For each event, public evidence was reconstructed approximately 30 days prior to failure and scored retroactively under the ACI framework. Criteria that demonstrated strong discriminative power were assigned higher weights; criteria that failed the discrimination test were reduced. Thresholds are anchored to empirical boundary values observed in the failure dataset and to established regulatory and industry precedents. The result is evidence-based, not narrative-based. Historical reconstruction testing across the observed failure dataset identified all reconstructed failures as ELEVATED or HIGH at T-30. Results remain subject to dataset limitations and ongoing recalibration.

5.2 Two Phases, Eight Modules, One Universal Architecture

Module	Product Category	Primary Criterion (Weight)
M1A — BTC Collateral Lending	Custodial and non-custodial BTC-collateralised lending	Collateral Control (35%)
M1B — Treasury Preferred Shares	BTC-treasury preferred-share instruments	BTC Coverage Ratio / SDACR (30%)
M1C-A — CeFi Stablecoins	Centralised stablecoin yield platforms	Solvency Verification (35%)
M1C-B — DeFi Stablecoins	On-chain stablecoin yield protocols	Protocol Security (40%)

M4 — Infrastructure	Bitcoin mining and validator operators	Unit Economics category (30%)
M5 — Market Neutral	Basis, funding-rate, hybrid market-neutral	Strategy Robustness category (30%)
M6 — Venture	Digital asset venture allocations and token launches	Team Quality category (30%)
M7 — Tokenised RWA	On-chain real-world assets	Asset Quality category (30%)
M8 — Volatility	Options, structured products, DeFi option vaults	Counterparty Quality (15–30% segment-adjusted)

5.3 Universal Scoring Formula

```
Raw Score   =  Σ ( criterion_score_i × weight_i )    for i across all module criteria
Adjusted Score  =  Raw Score  +  CascadePenalty(criterion_scores)
Final Score  =  clamp( round( Adjusted Score × Duration_Multiplier ), 0, 100 )
```

Where: each criterion_score_i is in [0, 100]; Σ weight_i = 1.00 per module; CascadePenalty returns −5 when three or more criterion_score_i values fall below 40, otherwise 0; Duration_Multiplier varies by module per §5.4. Hard caps (§7) may apply post-formula and override band assignment.

5.4 Duration Multiplier Schedule

Longer commitment durations carry compounding structural risk. The duration multiplier scales the risk score only — it never affects projected income. M5 (Market Neutral) and M6 (Venture) short-circuit the multiplier to 1.000 because their underlying instruments have no fixed maturity. M7 (Tokenised RWA) applies a subtype-driven multiplier instead.

Duration (Months)	Multiplier · Typical Context
≤ 3	1.000 — Quarterly rolling positions
≤ 6	1.050 — Semi-annual horizons
≤ 12	1.100 — Annual lock-ups
≤ 24	1.175 — Multi-year commitments
> 24	1.250 — Long-duration ventures

M7 subtype multipliers (planner-applied): tokenized_tbill 0.95, private_credit 1.10, real_estate 1.15, structured_product 1.05.

5.5 Linear Interpolation Protocol

For continuous-value criteria (HV30, HV1Y, depeg, LTV, BTC coverage, TVL), ACI applies linear interpolation between published anchor pairs. Score transitions are smooth and proportional. Categorical criteria retain step-function scoring.

```
Interpolated_Score = low_score + ( (V − lower) / (upper − lower) ) × ( high_score − low_score )
```

Scores are rounded to integer before use in the weighted sum. Below the lowest anchor, the lowest anchor's score applies. Above the highest anchor, the engine applies the configured aboveLastAnchor policy: "last_score" preserves the last anchor's score (descending tables); "cap" returns 100 (ascending tables).

6 METHODOLOGY INTEGRITY

6.1 Three Operating Principles

Principle 1 — Determinism

Every risk indicator is computed by a deterministic algorithm. Given the same publicly available inputs and the published weights, any analyst can independently replicate any ACI output. Determinism is enforced at the code level. The platform’s continuous integration pipeline runs SHA-256 determinism tests covering all currently active scoring modules; a determinism regression is a critical control exception that prevents release.

Principle 2 — Full Transparency

Every weight, sub-criterion, formula, data source, and hard rule is published in this document. No inputs are confidential. Criterion-level score breakdowns are surfaced in the platform's score-detail panel for every scenario.

Principle 3 — User Constructs, ACI Computes

Every output is traceable to user inputs. The user is responsible for the inputs. ACI is responsible for the computation being deterministic and methodology-consistent. This separation supports ACI's classification as an analytics service.

6.2 Score Versioning Policy

Score versioning is append-only. Every score references the framework version under which it was computed. Scores produced under a prior version remain valid and permanently retrievable. No retroactive revision. Version changes require formal CRO sign-off plus 30-day subscriber notification before taking effect.

6.3 Worst-Case Protocol

> **Worst-Case Protocol**
>
> When evidence quality for a criterion is insufficient or contradictory — confidence rating below 0.70 or conflicting evidence detected — the engine substitutes a CRO-approved conservative worst-case score for that criterion. Undisclosed information is always treated as the most adverse available outcome.
>
> Substitution is automatic and cannot be bypassed at the analyst level. Any exception requires CRO sign-off and a permanent audit log entry.

7 STRUCTURAL HARD RULES

7.1 Hard Rules Registry

Rule	Trigger	Effect
ALGO_STABLECOIN_CAP	Stablecoin peg_type = Algorithmic	Final score capped at 20
HIGH_DEPEG_OVERRIDE	Depeg > 1.5% over 90-day window	Forced HIGH band
HV30 cap	HV30 > 0.35	Market risk criterion capped at 5
CASCADE_PENALTY_APPLIED	three or more criteria scoring below 40	Minus 5 to raw score
VOLATILITY_REGIME_ACTIVE	HV30 / HV1Y regime multiple > 1.5×	Market risk shifts to HV30-dominant
REGULATORY_CONTROVERSY_MODIFIER	Provider under regulatory inquiry	Modifier applied to jurisdictional sub-score
ALL_CRITERIA_NOT_APPLICABLE	All criteria flagged not_applicable	Raw score → 0
VOLATILITY_NEGATIVE_CONVEXITY_LARGE_CAPITAL_BAND	M8 NEGATIVE convexity at large capital	Forced ELEVATED band floor
MN_LEVERAGE_NO_STOP_CAP	M5 leverage > 5× without stop-loss	Final score capped at 30
MN_UNKNOWN_COUNTERPARTY_CAP	M5 unknown counterparty	Final score capped at 30
MN_NO_POSITION_VISIBILITY_CAP	M5 no position visibility	Final score capped at 30
VENTURE_ANON_TEAM_CAP	M6 anonymous team	Final score capped at 30
RWA_OFFSHORE_NO_AUDIT_CAP	M7 offshore + no audit	Final score capped at 30
RWA_SELF_CUSTODY_CAP	M7 self-custody no segregation	Final score capped at 30

7.2 Empirical Rationale

Algorithmic stablecoin cap

The TerraUSD collapse in May 2022 — and prior failures of Iron Finance and USDN — confirmed that algorithmic peg mechanisms exhibit catastrophic non-linear failure modes. 100% of algorithmic stablecoins in the failure dataset experienced complete peg collapse within 18 months of weighted scoring above 60.

HIGH_DEPEG_OVERRIDE

The 1.5% threshold is empirically anchored. Historical analysis of stablecoin stress events shows deviations beyond 1.5% exhibit a materially different recovery profile. The USDD near-depeg of 1.9% (mid-2022) and the USDC SVB-related depeg of approximately 10% (March 2023) both occurred above this threshold and were associated with significant structural stress.

HV30 cap

HV30 ≥ 0.35 indicates regime-shift volatility consistent with crisis conditions. The hard cap is set at 5 to remove a step discontinuity at the boundary.

Cascade Penalty

Reflects the empirical observation that simultaneous deterioration across multiple criteria signals systemic stress that a simple weighted average underestimates. Celsius, BlockFi, and Voyager each failed with three or more criteria simultaneously below 40 in the months preceding restructuring.

8 SCORE BAND FRAMEWORK

Every ACI Risk Indicator falls into one of four bands. Band assignment is deterministic — no analyst discretion applies. Hard caps may force a band reclassification regardless of composite score.

Band · Score Range	Platform Behaviour
LOW · 80–100	Eligible for inclusion in Reference Scenario allocations. No allocation cap. Highest-quality providers within the universe.
MEDIUM · 60–79	Eligible with risk-disclosure note displayed prominently. Solid providers with manageable risk factors.
ELEVATED · 40–59	User-defined scenario only. Maximum 30% allocation cap applied. Excluded from default Reference Scenario auto-selection. Material risk factors present requiring explicit consideration.
HIGH · 0–39	Excluded from Reference Scenario. Available in user-defined scenarios with mandatory risk disclosure acknowledgment. Highest structural risk category.

9 CROSS-CUTTING ENGINE MECHANICS

9.1 Bucket Score Map

The Bucket Score Map converts categorical bucket values into 0–100 numeric scores. Each module has its own bucket map covering all of its criteria. The active map is the v1.0 BSM, embedded under ACI Framework v1.0. Modifications require written CRO sign-off plus 30-day subscriber notification per §14.2.

9.2 Convexity Engine

Each module attaches a convexity classifier — POSITIVE, NEUTRAL, or NEGATIVE — describing the structural payoff shape under stress. Display-only — does not affect score, except where the M8

NEGATIVE convexity at large capital triggers VOLATILITY_NEGATIVE_CONVEXITY_LARGE_CAPITAL_BAND.

Module	Default Convexity	Rationale
M1A	NEUTRAL	Symmetric exposure to BTC + counterparty
M1B	NEUTRAL	Fixed-coupon preferred backed by BTC reserves
M1C-A	NEUTRAL	Fixed-yield CeFi stablecoin lending
M1C-B	NEUTRAL	Smart-contract yield with peg-stability tail
M4	NEUTRAL	BTC-revenue exposure with energy-cost floor
M5	NEUTRAL	Funding-rate harvest with execution-quality tail
M6	POSITIVE	Power-law upside; the platform's convexity anchor
M7	NEUTRAL	Underlying-credit exposure with redemption tail
M8	NEGATIVE	Short-vol positions; explicit tail-risk surfacing

9.3 TVL Quality Discount (DeFi)

DeFi TVL is discounted before scoring to account for non-sticky, incentive-driven liquidity.

TVL Quality Tier	Discount	Definition
organic_sticky	0%	Capital persists without incentives
moderate_incentivised	10%	Modest farming incentives in place
heavily_incentivised	30%	Substantial farming incentives drive deposits
mercenary_dominated	50%	Capital follows the highest-yield farm; high churn

10 MODULE METHODOLOGY

Each module is documented criterion-by-criterion with the full bucket score table for every criterion. The complete BSM source path is cited in each section for code-level reproducibility.

10.1 Module M1A — BTC Collateral Lending

Scope. Custodial and non-custodial credit facilities collateralised by BTC. The user pledges BTC; the platform issues credit (USD or stablecoin) or earns yield on the collateral. ACI scores the platform's custody discipline, liquidation mechanics, jurisdictional accountability, and operating track record. BTC price risk is shown separately on the platform and is not folded into the structural score.

C1 — Transparency and Disclosure · Weight: 20%

Measures whether the platform gives users the information necessary to understand the safety conditions around their pledged BTC. Primary driver: Proof of Reserves (PoR) frequency. Empirical analysis of CeFi failures showed disclosure quality alone has limited discriminative power — several platforms with active PoR attestations subsequently collapsed due to undisclosed structural risks.

BUCKET VALUE	SCORE
Monthly attestation by named independent auditor	**100**

On-chain individual vault verification (Unchained 2-of-3 multisig)	**90**
Quarterly attestation	**75**
Annual attestation	**50**
No proof of reserves	**0**

C2 — Collateral Control and Liquidation Mechanics · Weight: 35% (composite)

Custody Model is the dominant single predictor of platform failure outcomes in the empirical dataset: every major CeFi platform that collapsed with material client losses operated commingled or undisclosed custody models.

```
Collateral_Control_Score = round( CustodyModel × 0.70 + CollateralTopUp × 0.30 )
```

C2.1 — Custody Model · Weight: 70% of C2

BUCKET VALUE	SCORE
Fully segregated custody, disclosed to clients	**100**
Pooled custody, disclosed to clients	**60**
Commingled with operational funds	**10**
Unknown / not disclosed	**0**

C2.2 — Collateral Top-Up Speed · Weight: 30% of C2

BUCKET VALUE	SCORE
Instant top-up required upon margin call	**100**
Same business day	**75**
Delayed (2–5 business days)	**30**
No top-up mechanism — immediate liquidation triggered	**0**

C2.3 — Liquidation LTV · Linear interpolation

BUCKET VALUE	SCORE
≤ 75%	**100**
76–80%	**80**
81–85%	**60**
86–90%	**30**
> 90%	**0**

C2.4 — Maximum LTV Offered · Linear interpolation

BUCKET VALUE	SCORE
≤ 40%	**100**
41–50%	**85**

51–60%	**65**
61–70%	**40**
71–80%	**15**
> 80%	**0**

C3 — Jurisdiction and Legal Enforceability · Weight: 15%

Categorical step-function scoring across four tiers reflecting the regulatory framework available to clients.

BUCKET VALUE	SCORE
Tier 1 — CH, US, UK, SG, EU (MiCA). Established regulatory frameworks.	**100**
Tier 2 — Gibraltar, HK, Cayman Islands. Frameworks exist; protections narrower.	**65**
Tier 3 — Malta, BVI, Seychelles. Lighter-touch regulation.	**25**
Tier 4 — Offshore, anonymous, unregistered.	**0**

C4 — Structural and Custody Risk (Rehypothecation) · Weight: 25%

Empirical analysis identifies rehypothecation status as the second-strongest predictor of platform failure. Undisclosed rehypothecation was the central structural mechanism in the Celsius collapse ($4.7B in client losses) and was present in every major CeFi lending failure between 2020 and 2023.

BUCKET VALUE	SCORE
No rehypothecation — client BTC ring-fenced	**100**
Rehypothecation disclosed in Terms of Service	**25**
Rehypothecation not disclosed (undisclosed)	**0**

C5 — Track Record and Operational Resilience · Weight: 5%

BUCKET VALUE	SCORE
Established & Regulated — 5+ years, full regulatory licence, survived market cycle	**100**
Mature & Licensed — 3–5 years, holds regulatory licence or registration	**75**
Operational & Registered — 1–3 years, basic registration	**50**
New or Unproven — < 1 year operating, OR no regulatory history	**15**

Custodial: AMINA Bank (CH, T1), Sygnum Bank (CH, T1), Xapo (GI, T2), Matrixport (SG, T2), Ledn (KY, T3), Nexo (BG, T2), Salt Lending (US, T1), Strike (US, T1), Abra (US, T1)

Non-custodial: Unchained (US, T1), Hodl Hodl P2P (MT, T2), Figure (US, T1)

10.2 Module M1B — Treasury Preferred Shares

Scope. BTC-treasury preferred share instruments — equity instruments issued by companies whose primary corporate treasury asset is Bitcoin, where preferred dividends are economically backed by that BTC holding. BTC price risk is shown separately on the platform.

C1 — BTC Coverage Ratio / SDACR · Weight: 30%

BTC-treasury preferred shares are equity instruments — not debt. In a wind-down scenario, senior secured debt, senior unsecured debt, and accounts payable rank ahead of preferred equity in the capital-structure waterfall. The Senior-Debt-Adjusted Coverage Ratio (SDACR) correctly models the preferred shareholder's economic position.

```
SDACR = ( BTC_Holdings × Current_BTC_Price − Total_Senior_Debt ) ÷
Total_Preferred_Obligations
```

BUCKET VALUE	SCORE
Extreme — SDACR > 20× preferred obligations	**100**
Strong — SDACR 10–20×	**80**
Adequate — SDACR 5–10×	**60**
Thin — SDACR 1.5–5×	**30**
At Risk — SDACR < 1.5× OR undisclosed (hard floor)	**0**

C2 — Income Mechanism · Weight: 25%

The contractual structure of dividend payment determines the legal protection available to preferred holders.

BUCKET VALUE	SCORE
Fixed Contractual — fixed rate, legally binding; non-payment triggers default	**100**
Fixed Board-Declared — fixed rate but board can suspend without legal default	**60**
Fixed FX — fixed rate in foreign currency; USD investors carry FX exposure	**50**
Variable Formula — variable rate tied to a defined, published formula	**40**
Fully Discretionary — entirely at board discretion	**10**

C3 — Market Risk Composite · Weight: 20% (composite)

Composite of three sub-components capturing market risk in distinct dimensions. Volatility regime detection: HV30 / HV1Y > 1.5× shifts the composite to HV30-dominant weighting (HV30 70%, HV1Y 30%).

```
Market_Risk_Score = round( HV1Y × 0.50 + Price_to_Par × 0.30 + Liquidity × 0.20 )
```

C3.1 — HV1Y Tier · Weight: 50% of C3

BUCKET VALUE	SCORE
Very Low — < 15%	**100**
Low — 15–25%	**80**
Moderate — 25–40%	**55**
High — 40–60%	**25**
Extreme — > 60%	**0**

C3.2 — Price-to-Par Tier · Weight: 30% of C3

BUCKET VALUE	SCORE
At Par — 98–102%	**100**
Near Par — 90–98%	**80**
Moderate Discount — 75–90%	**50**
Deep Discount — 60–75%	**20**
Distressed — below 60%	**0**

C3.3 — Liquidity Tier (30-day average volume) · Weight: 20% of C3

BUCKET VALUE	SCORE
Institutional — > $100M	**100**
Liquid — $20M – $100M	**75**
Moderate — $5M – $20M	**50**
Thin — $1M – $5M	**20**
Illiquid — < $1M	**0**

C4 — Convertibility Risk · Weight: 17%

BUCKET VALUE	SCORE
Non-Convertible — preferred status permanent	**100**
Holder-Optional Convertible — at holder discretion only	**60**
Issuer-Forced / Trigger-Based — issuer can force OR auto-trigger under adversity	**20**

C5 — Issuer Maturity · Weight: 8%

BUCKET VALUE	SCORE
Institutional & Established — 3+ years BTC treasury, SEC-registered, survived full cycle	**100**
Listed & Emerging — listed but < 3 years BTC treasury history	**65**
Private or New — private entity or BTC treasury < 12 months	**25**
Undisclosed — BTC treasury cannot be independently verified	**0**

Instruments: STRF (Strategy 10% perpetual), STRK (convertible Series A), STRC (cumulative), STRD (Strive discretionary), SATA (Strive Series A), STRE (EUR-denominated)

10.3 Module M1C-A — CeFi Stablecoins

Scope. Centralised stablecoin yield platforms. Most platforms are stablecoin distributors, not issuers. M1C-A does not assess stablecoin issuer reserve quality — it assesses the platform's own financial safety. Core question: if this platform fails, do clients recover their stablecoins?

C1 — Solvency Verification · Weight: 35%

Solvency Verification is the dominant criterion because empirical analysis confirms it as the strongest single predictor of CeFi platform failure.

BUCKET VALUE	SCORE
Big 4 audited annual — accounts audited by a Big 4 firm, filed with SEC/FCA/FINMA	**100**
Independent audited annual — audited by recognised CPA, publicly filed	**85**
Proof of Reserves (PoR) quarterly — by named CPA	**70**
Proof of Reserves (PoR) annual — by named CPA	**55**
Self-reported only — platform publishes own figures; no named auditor	**20**
No disclosure — no independent verification	**0**

C2 — Regulatory Accountability · Weight: 20% (composite)

Three sub-questions, all verifiable from public registers.

```
Reg_Accountability_Score = round( Q1 × 0.40 + Q2 × 0.40 + Q3 × 0.20 )
```

C2.Q1 — Is the yield product within the platform's licence? · 40% of C2

BUCKET VALUE	SCORE
Explicitly licensed	**100**
Probably covered	**70**
Unclear	**40**
Exchange only	**20**
Unlicensed	**0**

C2.Q2 — Does the regulator have enforcement powers? · 40% of C2

BUCKET VALUE	SCORE
Prudential supervisor	**100**
Conduct regulator	**80**
AML only	**40**
Registration only	**10**

C2.Q3 — Client compensation or complaints mechanism? · 20% of C2

BUCKET VALUE	SCORE
Statutory compensation scheme (FSCS, SIPC)	**100**
Regulatory complaints with binding arbitration	**70**
Voluntary scheme or ombudsman only	**40**
None	**0**

C3 — Yield Commitment · Weight: 10%

Yield Commitment is weighted at 10% — the lowest non-trivial weight in M1C-A — because empirical analysis showed near-zero discriminative power: every major CeFi platform that collapsed maintained attractive promoted yield offers in the months preceding failure.

BUCKET VALUE	SCORE
Contractual Fixed — fixed rate, legally binding	**100**
Contractual Variable — variable rate tied to a defined formula	**75**
Disclosed Discretionary — platform may change rate with notice	**45**
Promotional Disclosed — promotional rate; no contractual floor	**20**
Promotional Undisclosed — rate may be removed without notice	**0**

C4 — Liquidity · Weight: 25% (composite)

Liquidity is weighted at 25% because TVL drawdown is a leading indicator of CeFi platform failure: TVL declines preceded failures by a median of approximately 47 days in the historical dataset.

```
Liquidity_Score = round( TVL × 0.70 + Withdrawal_Speed × 0.30 )
```

C4.1 — TVL Tier (CeFi) · 70% of C4

BUCKET VALUE	SCORE
Above $10B	**100**
$1B – $10B	**85**
$100M – $1B	**65**
$10M – $100M	**40**
Below $10M	**10**

C4.2 — Withdrawal Speed · 30% of C4

BUCKET VALUE	SCORE
Instant — immediate redemption	**100**
< 7 days	**65**
< 30 days	**30**
Locked — no redemption path	**0**

C5 — Jurisdiction · Weight: 10% (composite)

```
Jurisdiction_Score = round( Incorporation × 0.30 + Product_Oversight × 0.70 )
```

C5.1 — Incorporation Tier · 30% of C5

BUCKET VALUE	SCORE
Tier 1 G20 — UK, US, EU, SG, CH	**100**

Tier 2 Established	**75**
Tier 3 Light-Touch	**40**
Tier 4 Opaque	**0**

C5.2 — Product Oversight Tier · 70% of C5

BUCKET VALUE	SCORE
Prudential licensed — full balance-sheet oversight	**100**
VASP full Tier 1	**80**
VASP full Tier 2	**60**
Registered but not licensed	**25**
Unregulated	**0**

Representative Providers: Coinbase (US, T1), Kraken Earn (US, T1), Nexo (BG, T2), Bybit (UAE, T3), OKX Earn (SC, T3), Binance FZE (UAE, T4), YouHodler (CH, T2), Matrixport (SG, T2), Ledn (KY, T3)

10.4 Module M1C-B — DeFi Stablecoins

Scope. DeFi stablecoin yield protocols. DeFi protocols have no company balance sheet, no traditional custody risk, and no regulatory licence. The risk is code risk and governance risk. Core question: is the code secure, and can governance actors attack users?

C1 — Protocol Security · Weight: 40% (composite)

Protocol Security carries the highest weight in M1C-B because empirical analysis of DeFi protocol exploits (2020–2025) confirms audit depth and battle-test duration as the two most predictive criteria for protocol survival.

```
Protocol_Security_Score = round( Audit_Depth × 0.50 + Battle_Test × 0.50 )
```

C1.1 — Audit Depth Tier · 50% of C1

BUCKET VALUE	SCORE
Tier 1 multi-audit — 3+ Tier-1 firms + Immunefi bug bounty > $1M	**100**
Tier 1 dual audit — 2 Tier-1 firms + active bug bounty	**80**
Recognised single audit — 1 recognised audit + bug bounty	**55**
Unrecognised audit — non-Tier-1 firm	**25**
Unaudited — no published security audit	**0**

C1.2 — Battle-Test Tier · 50% of C1

BUCKET VALUE	SCORE

Sustained 3yr no exploit — current codebase > 3 years without material exploit	**100**
Sustained 2yr no exploit — 2–3 years without exploit	**80**
Sustained 1yr no exploit — 1–2 years without exploit	**60**
Under 1yr — less than 1 year of operation	**35**
Post-exploit recovered — prior exploit; full user reimbursement	**25**
Active exploit history — material exploit, incomplete or no recovery	**0**

C2 — Governance Risk · Weight: 30% (composite)

Governance Risk reflects the empirical finding that governance attack is a major DeFi failure vector. The Beanstalk ($182M) and Mango Markets ($117M) exploits were governance attacks, not code exploits.

```
Governance_Risk_Score = round( Immutability × 0.50 + Regulatory_Record × 0.50 )
```

C2.1 — Immutability Tier · 50% of C2

BUCKET VALUE	SCORE
Fully immutable — no upgrade path, no proxy, no admin key	**100**
Timelock 72h+ — governance vote + 72h+ timelock	**80**
Timelock 24–72h — governance vote + 24–72h timelock	**60**
Governance no timelock — upgradeable by governance	**30**
Admin key controlled — team holds upgrade keys	**10**

C2.2 — Regulatory Record Tier · 50% of C2

BUCKET VALUE	SCORE
No action — geofenced — restricted-jurisdiction access blocked	**100**
No action — open — operates in ungeofenced grey zones	**75**
Under inquiry — formal action possible	**40**
Formal action — subject to regulatory enforcement	**0**

C3 — Yield Transparency · Weight: 15%

BUCKET VALUE	SCORE
Fully on-chain documented — all yield parameters on-chain with documentation	**100**
On-chain underdocumented — parameters on-chain, documentation incomplete	**75**
Partially on-chain — some parameters on-chain; others off-chain	**45**

Off-chain disclosed — yield determined off-chain with publicly disclosed methodology	**20**
Not disclosed — yield methodology not disclosed or verifiable	**0**

C4 — Liquidity · Weight: 10% (composite)

Same composite structure as M1C-A: TVL × 0.70 + Withdrawal_Speed × 0.30. DeFi TVL anchors are larger to reflect the higher capital scale of mature DeFi protocols. DeFi TVL is discounted by quality tier (§9.3) before scoring.

BUCKET VALUE	SCORE
Above $5B	**100**
$1B – $5B	**85**
$500M – $1B	**70**
$200M – $500M	**50**
$50M – $200M	**25**
Below $50M	**0**

C5 — Peg Stability · Weight: 5%

Scored from the maximum peg deviation over 90 days. 5% weight but triggers HIGH_DEPEG_OVERRIDE if 90-day max exceeds 1.5%.

BUCKET VALUE	SCORE
≤ 0.10% (10 bps)	**100**
0.10–0.30% (10–30 bps)	**80**
0.30–0.50% (30–50 bps)	**60**
0.50–1.50% (50–150 bps)	**20**
> 1.50% — HIGH_DEPEG_OVERRIDE triggered → forced HIGH band	**0**

Representative Providers: Aave v3, Compound v3, Spark Protocol, Morpho Blue, Ethena USDe

10.5 Module M4 — Infrastructure

Scope. Bitcoin mining and validator operators. Yield from block rewards, transaction fees, or staking. ACI scores unit economics under BTC price shocks, energy structure, operational reliability, counterparty discipline, transparency, and capacity. Most M4 criteria use a standard four-tier bucket map (Strong / Adequate / Weak / Inadequate).

Standard M4 Bucket Map	Score
Strong	90

Adequate	70
Weak	45
Inadequate	20

Unit Economics — 30% (3 criteria)

C1 — Breakeven BTC Price · Weight: 12%

Distance to monthly cashflow zero. Strong tier indicates substantial buffer above current BTC price.

BUCKET VALUE	SCORE
Strong (substantial buffer above current BTC price)	**90**
Adequate	**70**
Weak	**45**
Inadequate (cashflow negative below current BTC)	**20**

C2 — Energy Cost per kWh · Weight: 10%

BUCKET VALUE	SCORE
Strong — < $0.035/kWh	**90**
Adequate — $0.035–$0.055/kWh	**70**
Weak — $0.055–$0.075/kWh	**45**
Inadequate — > $0.075/kWh	**20**

C3 — Margin Stability · Weight: 8%

Net-margin variability across recent quarters. Uses standard M4 bucket map.

Reliability — 25% (3 criteria)

C4 — Uptime Percentage · Weight: 10%

BUCKET VALUE	SCORE
Strong — > 99%	**90**
Adequate — 97–99%	**70**
Weak — 95–97%	**45**
Inadequate — < 95%	**20**

C5 — Historical Downtime · Weight: 8%

Frequency and severity of operational incidents. Uses standard M4 bucket map.

C6 — Maintenance Quality · Weight: 7%

Maintenance cadence and disclosure. Uses standard M4 bucket map.

Counterparty — 20% (3 criteria)

C7 — Legal Structure · Weight: 7%

Bankruptcy-remoteness and corporate structure. Uses standard M4 bucket map.

C8 — Asset Custody · Weight: 7%

Audited qualified custody is the only tier that fully survives provider insolvency.

BUCKET VALUE	SCORE
Qualified Custodian (Audited)	**95**
Qualified Custodian (Unaudited)	**75**
Self-Custody (Audited)	**55**
Self-Custody (Self-Reported)	**25**

C9 — Jurisdiction · Weight: 6%

BUCKET VALUE	SCORE
Tier 1 — Transparent (US, CH, UK, SG, EU)	**95**
Tier 2 — Regulated	**75**
Tier 3 — Emerging	**50**
Offshore Opaque	**20**

Transparency — 15% (2 criteria)

C10 — Real-Time Dashboard · Weight: 8%

Live operational data disclosure. Uses standard M4 bucket map.

C11 — Audit Frequency · Weight: 7%

External audit cadence. Uses standard M4 bucket map.

Scalability — 10% (1 criterion)

C12 — Capacity Scaling · Weight: 10%

Demonstrated growth without operational degradation. Uses standard M4 bucket map.

> **Representative Providers:** Marathon Digital, Riot Platforms, Cipher Mining, Hut 8, Core Scientific, Compass Mining

10.6 Module M5 — Market Neutral

Scope. Basis, funding-rate, and hybrid market-neutral strategies. Yield from harvesting persistent funding-rate or basis-trade premiums. Continuously rebalanced — duration multiplier short-circuits to 1.000.

Strategy Robustness — 30% (3 criteria)

C1 — Strategy Type · Weight: 10%

BUCKET VALUE	SCORE
funding_arb_documented	**100**
basis_trading_documented	**85**
delta_neutral_documented	**75**
hybrid_undisclosed	**50**
strategy_undisclosed	**20**

C2 — Sharpe Ratio · Weight: 10%

BUCKET VALUE	SCORE
sharpe_above_two	**100**
sharpe_one_to_two	**75**
sharpe_below_one	**40**
sharpe_not_disclosed	**15**

C3 — AUM Disclosure · Weight: 10%

BUCKET VALUE	SCORE
aum_audited_disclosed	**100**
aum_self_reported	**65**
aum_range_only	**40**
aum_not_disclosed	**10**

Risk Management — 25% (3 criteria)

C4 — Leverage · Weight: 9%

Hard cap: leverage > 5× without stop-loss caps the final score at 30 (MN_LEVERAGE_NO_STOP_CAP).

BUCKET VALUE	SCORE
leverage_below_two	**100**
leverage_two_to_three	**75**
leverage_three_to_five	**45**
leverage_above_five_no_stop	**0**

C5 — Counterparty · Weight: 8%

Hard cap: unknown counterparty caps the final score at 30 (MN_UNKNOWN_COUNTERPARTY_CAP).

BUCKET VALUE	SCORE
tier_one_prime_segregated	**100**
tier_one_prime_commingled	**75**
tier_two_exchange_only	**50**

unknown_counterparty	**0**

C6 — Collateral Segregation · Weight: 8%

BUCKET VALUE	SCORE
fully_segregated_audited	**100**
segregated_self_reported	**70**
partial_segregation	**40**
commingled_no_segregation	**10**

Execution Quality — 20% (3 criteria)

C7 — Funding Rate Environment · Weight: 7%

BUCKET VALUE	SCORE
consistently_positive_90d	**100**
mostly_positive_with_dips	**75**
volatile_funding	**40**
negative_funding_environment	**15**

C8 — Venue · Weight: 7%

BUCKET VALUE	SCORE
tier_one_cme_regulated	**100**
tier_one_deribit_binance	**85**
multi_venue_mixed_tier	**60**
tier_two_only	**30**

C9 — Pricing Infrastructure · Weight: 6%

BUCKET VALUE	SCORE
institutional_grade_verified	**100**
proprietary_documented	**75**
vendor_standard	**50**
pricing_undisclosed	**20**

Liquidity Access — 15% (2 criteria)

C10 — Exchange Diversity · Weight: 8%

BUCKET VALUE	SCORE
five_plus_tier_one	**100**

three_to_four_tier_one	**80**
one_to_two_tier_one	**50**
single_exchange_concentrated	**20**

C11 — Position Visibility · Weight: 7%

Hard cap: no position visibility caps the final score at 30 (MN_NO_POSITION_VISIBILITY_CAP).

BUCKET VALUE	SCORE
real_time_full_disclosure	**100**
daily_reporting	**70**
monthly_reporting	**40**
no_position_visibility	**0**

Transparency — 10% (4 criteria)

C12 — Track Record Duration · Weight: 4%

BUCKET VALUE	SCORE
three_plus_years_verified	**100**
one_to_three_years	**70**
under_one_year	**35**
no_track_record	**10**

C13 — Drawdown History · Weight: 3%

BUCKET VALUE	SCORE
max_drawdown_below_five	**100**
drawdown_five_to_fifteen	**70**
drawdown_above_fifteen	**35**
drawdown_not_disclosed	**15**

C14 — Strategy Disclosure · Weight: 2%

BUCKET VALUE	SCORE
full_prospectus_published	**100**
summary_disclosed	**60**
verbal_only	**30**
not_disclosed	**10**

C15 — Audit · Weight: 1%

BUCKET VALUE	SCORE
big_four_annual	**100**
reputable_firm_annual	**80**

audit_exists_older	**50**
no_audit	**10**

Representative Providers: Wintermute, GSR, Cumberland, Amber Group, B2C2

10.7 Module M6 — Venture

Scope. Digital asset venture allocations — early-stage token launches and pre-listing positions. Power-law return distributions. Convexity classifier defaults to POSITIVE — M6 is the platform's convexity anchor. Duration multiplier short-circuits to 1.000.

Team Quality — 30% (3 criteria)

C1 — Track Record · Weight: 12%

BUCKET VALUE	SCORE
multiple_exits_top_decile	**100**
exits_verified_solid	**80**
early_track_record_promising	**55**
no_verifiable_exits	**20**

C2 — Execution Capability · Weight: 10%

BUCKET VALUE	SCORE
institutional_operations_soc2	**100**
structured_team_processes	**75**
lean_team_early_stage	**45**
founder_only_undocumented	**15**

C3 — Team Depth · Weight: 8%

Hard cap: anonymous team caps the final score at 30 (VENTURE_ANON_TEAM_CAP).

BUCKET VALUE	SCORE
deep_team_three_plus_gp	**100**
core_team_two_gp	**70**
key_person_risk_one_gp	**35**
anonymous_team	**0**

Tokenomics Expertise — 20% (3 criteria)

C4 — Supply Schedule · Weight: 7%

BUCKET VALUE	SCORE

transparent_schedule_audited	**100**
schedule_disclosed_not_audited	**70**
partial_disclosure	**40**
schedule_not_disclosed	**10**

C5 — Unlock Risk · Weight: 7%

BUCKET VALUE	SCORE
long_vesting_cliff_two_plus	**100**
standard_vesting_one_year	**70**
short_vesting_below_one_year	**35**
no_vesting_immediate_unlock	**5**

C6 — FDV Reasonableness · Weight: 6%

BUCKET VALUE	SCORE
fdv_below_ten_x_revenue	**100**
fdv_ten_to_fifty_x	**65**
fdv_above_fifty_x	**25**
fdv_not_disclosed	**10**

Market Potential — 20% (3 criteria)

C7 — TAM · Weight: 7%

BUCKET VALUE	SCORE
tam_above_ten_billion	**100**
tam_one_to_ten_billion	**70**
tam_below_one_billion	**35**
tam_not_quantified	**15**

C8 — Narrative Strength · Weight: 7%

BUCKET VALUE	SCORE
category_defining_narrative	**100**
strong_differentiated_thesis	**75**
competitive_crowded_space	**40**
narrative_unclear	**15**

C9 — Category Maturity · Weight: 6%

BUCKET VALUE	SCORE
established_category_growing	**100**

emerging_category_validated	**70**
nascent_high_uncertainty	**35**
experimental_unproven	**10**

Investor Base — 15% (2 criteria)

C10 — Smart Money · Weight: 9%

BUCKET VALUE	SCORE
tier_one_multi_lead	**100**
tier_one_lead_only	**75**
tier_two_mixed	**40**
no_smart_money	**0**

C11 — Syndicate Diversity · Weight: 6%

BUCKET VALUE	SCORE
diverse_strategic_plus_financial	**100**
primarily_financial_investors	**65**
single_lead_concentrated	**35**
syndicate_not_disclosed	**15**

Liquidity Path — 15% (3 criteria)

C12 — Listing Likelihood · Weight: 6%

BUCKET VALUE	SCORE
tier_one_exchange_committed	**100**
tier_one_exchange_probable	**75**
tier_two_exchange_likely	**40**
listing_unclear	**10**

C13 — Exit Scenario · Weight: 5%

BUCKET VALUE	SCORE
multiple_exit_paths_defined	**100**
token_exit_only_credible	**60**
exit_path_unclear	**25**
no_exit_visibility	**5**

C14 — Vesting · Weight: 4%

BUCKET VALUE	SCORE
investor_vesting_aligned	**100**

standard_investor_terms	**65**
short_investor_vesting	**30**
vesting_not_disclosed	**10**

Representative Providers: a16z crypto, Polychain, Pantera, Paradigm, Dragonfly

10.8 Module M7 — Tokenised RWA

Scope. On-chain real-world asset platforms — sovereign bills, corporate credit, private credit, structured products. Subtype-driven duration multiplier applied at session-composition time by the planner.

Asset Quality — 30% (3 criteria)

C1 — Credit Risk Tier · Weight: 12%

Four categorical levels reflecting credit quality of the underlying assets.

BUCKET VALUE	SCORE
Sovereign-low-risk tier	**95**
Mid-grade credit tier	**85**
Sub-grade credit tier	**65**
High-risk unsecured tier	**40**

C2 — Duration Tier · Weight: 10%

BUCKET VALUE	SCORE
short_under_6m	**90**
medium_6m_to_18m	**65**
long_above_18m	**35**

C3 — Yield Sustainability · Weight: 8%

BUCKET VALUE	SCORE
overcollateralized	**100**
fully_collateralized	**85**
partially_collateralized	**55**
unsecured	**25**

Legal Structure — 25% (3 criteria)

C4 — Enforceability · Weight: 9%

BUCKET VALUE	SCORE
strong_enforceable	**90**
moderate_enforceable	**70**
weak_uncertain	**40**

C5 — Investor Protection · Weight: 9%

BUCKET VALUE	SCORE
strong_structural_protection	**90**
moderate_protection	**70**
weak_or_none	**40**

C6 — Jurisdiction · Weight: 7%

Feeds RWA_OFFSHORE_NO_AUDIT_CAP hard cap (offshore + no audit caps final score at 30).

BUCKET VALUE	SCORE
tier_one_regulated	**95**
tier_two_regulated	**75**
offshore_unregulated	**40**

Custody — 20% (2 criteria)

C7 — Custodian · Weight: 11%

Feeds RWA_SELF_CUSTODY_CAP hard cap (self-custody no segregation caps final score at 30).

BUCKET VALUE	SCORE
qualified_custodian	**95**
segregated_accounts	**80**
self_custody_no_segregation	**30**

C8 — Bankruptcy Remoteness · Weight: 9%

BUCKET VALUE	SCORE
fully_remote_spv	**95**
partially_remote	**70**
not_remote	**40**

Liquidity — 15% (2 criteria)

C9 — Redemption Terms · Weight: 8%

BUCKET VALUE	SCORE
daily_liquidity	**95**
periodic_liquidity	**70**
locked_no_redemption	**40**

C10 — Secondary Market · Weight: 7%

BUCKET VALUE	SCORE
active_secondary_market	**90**

limited_secondary	**65**
no_secondary	**40**

Transparency — 10% (2 criteria)

C11 — Reporting · Weight: 5%

BUCKET VALUE	SCORE
real_time_transparency	**95**
periodic_reporting	**75**
limited_disclosure	**50**

C12 — Audit · Weight: 5%

Feeds RWA_OFFSHORE_NO_AUDIT_CAP hard cap.

BUCKET VALUE	SCORE
audited_big4	**95**
audited_third_party	**75**
no_audit	**30**

Representative Providers: Securitize, Backed Finance, Ondo Finance, Maple Finance, Centrifuge

10.9 Module M8 — Volatility

Scope. Options, structured products, and DeFi option vaults. M8 applies segment-adjusted weights — the importance of each criterion varies by venue type. Convexity-aware: NEGATIVE convexity surfaced explicitly.

Segment-Adjusted Weights

Criterion	Institutional OTC · Exchange Venue · On-Chain DOV
Strategy Design	25% · 20% · 20%
Counterparty Quality	30% · 15% · 25%
Execution Infrastructure	15% · 25% · 15%
Track Record	15% · 15% · 10%
Transparency	15% · 25% · 30%

C1 — Strategy Design · Weight: 20–25%

BUCKET VALUE	SCORE
undisclosed_strategy	**10**
naked_directional	**25**
basic_covered	**55**
defined_risk	**75**

sophisticated_multi_leg	**90**

C2 — Counterparty Quality · Weight: 15–30%

BUCKET VALUE	SCORE
unknown	**10**
unrated_offshore	**35**
reputable_otc	**60**
regulated_entity	**80**
tier_one_regulated	**95**

C3 — Execution Infrastructure · Weight: 15–25%

BUCKET VALUE	SCORE
manual_otc	**25**
basic_venue	**50**
deep_liquid_venue	**75**
institutional_rfq	**85**
multi_venue_smart_routing	**95**

C4 — Track Record · Weight: 10–15%

BUCKET VALUE	SCORE
no_history	**10**
under_one_year	**30**
one_to_three_years	**55**
three_to_five_years	**75**
over_five_years	**90**

C5 — Transparency · Weight: 15–30%

BUCKET VALUE	SCORE
opaque	**10**
periodic_disclosure	**40**
quarterly_greeks_published	**65**
monthly_full_positions	**80**
realtime_onchain	**95**

Segment routing via provider_subtype: institutional_otc / exchange_venue / on_chain_dov. Convexity classifier defaults to NEGATIVE for short-vol positions. Hard cap: VOLATILITY_NEGATIVE_CONVEXITY_LARGE_CAPITAL_BAND forces ELEVATED band floor for large-capital negative-convexity positions.

Representative Providers: Deribit, Wintermute (OTC), Ribbon Finance, Friktion, Opyn

11 STRESS ENGINES

Every ACI module ships a deterministic stress projection. Phase 1 modules use parametric curve analysis computed against published threshold tables. Phase 2 modules ship multi-scenario stress engines with full per-scenario output. Stress outputs are scenario-specific computations — they are not forecasts, predictions, or recommendations.

11.1 Phase 1 — Parametric Curve Stress

Phase 1 stress is implemented as a parametric curve analysis: the engine projects how the composite score and key sub-metrics evolve under a parameterised shock variable (typically BTC price drawdown or stablecoin peg deviation). The user moves a single slider; the engine recomputes outputs deterministically.

M1A — BTC Collateral Lending Stress Curve

Scenario	Driver	Projected Output
BTC drawdown −30%	BTC price shock	Projected margin-call probability, post-shock LTV, liquidation distance
BTC drawdown −50%	BTC price shock	Liquidation cascade — margin-call queue, available top-up window
Custody event	Custody-failure modelling	Recovery scenario based on custody tier and rehypothecation status
Jurisdictional shutdown	Regulator action on platform	Asset-recovery scenario by jurisdiction tier

M1B — Treasury Preferred Shares Stress Curve

Scenario	Driver	Projected Output
BTC drawdown −30% / −50% / −70%	BTC price shock	Projected SDACR, dividend-breakeven distance, projected band
HV1Y volatility regime	Volatility spike	Market risk composite shift; HV30-dominant if regime active
Dividend suspension	Board-declared income halt	Income-mechanism band reset
Convertibility trigger	Forced equity conversion event	Convertibility band → 20 (hard floor)

M1C — Stablecoin Yield Stress Curve (CeFi + DeFi)

Scenario	Driver	Projected Output
TVL run	Withdrawal acceleration	Liquidity tier degradation; withdrawal-speed band reset
Withdrawal halt	Operational freeze	Liquidity → 0; band override possible
Depeg event 0.5–1.5%	Peg deviation cascade	Peg-stability band reset
Depeg event >	Major peg break	HIGH_DEPEG_OVERRIDE fires — forced HIGH

1.5%		band
Governance attack (DeFi)	Admin-key or vote attack	Governance Risk band reset
Smart-contract exploit (DeFi)	Code-level vulnerability	Battle-Test tier reset

11.2 Phase 2 — Multi-Scenario Stress Engines

Module	Fixed Stress Grid	Primary Output
M4 — Infrastructure	Energy shock · Difficulty drift · Combined adverse · Halving compression	Monthly cashflow projection + failure-mode classifier
M5 — Market Neutral	Rate compression · Correlation shock · Extreme drawdown · Liquidity crunch	APY contraction + drawdown projection
M6 — Venture	Vintage concentration · Macro contraction · Correlation shock · Early-redemption stress	IRR distribution shift + outcome reweighting
M7 — Tokenised RWA	Spread widening · Sovereign stress · Liquidity trigger · Combined adverse	Net yield + redemption-gate probability
M8 — Volatility	Spot drawdown · Time decay · Vol regime shift · Combined adverse	Tail-loss projection + premium capture

11.3 Stress Output Discipline

Stress outputs are deterministic computations of the engine under specified scenario inputs. The IC committee reads the stress outputs and decides — ACI never states what action should follow from a stress result. The forbidden-verb list (REJECT, APPROVE, INVEST, AVOID, BUY, SELL, HOLD, RECOMMEND, ADVISE, CONVICTION, PREFERENCE) does not appear in any stress output surface.

12 PORTFOLIO AGGREGATION

Two or more saved scenarios from at least two distinct modules can be composed into a portfolio. The aggregate score is capital-weighted and duration-adjusted.

12.1 Capital-Weighted Score

```
Portfolio Score = Σ (score_i × capital_i × df(duration_i)) / Σ (capital_i × df(duration_i))
```

Where df(duration) is the duration factor table from §5.4. The aggregate score is clamped to [0, 100] and rounded to integer.

12.2 Cross-Module Concentration

Severity	Trigger	What Is Surfaced
CRITICAL	Single counterparty > 50% of total portfolio capital	Critical alert; concentration above institutional threshold
WARNING	Single counterparty > 30% and ≤ 50%	Warning alert; concentration above diversification threshold
INFO	Top 3 counterparties combined > 70%	Informational note on systemic concentration
CROSS-MODULE	Same counterparty appears across ≥ 2 modules	Combined exposure listed; severity scaled by aggregate %

12.3 Canonical Module Order

Portfolio surfaces and IC packages display modules in canonical order: 1A → 1B → 1C → 4 → 5 → 7 → 8 → 6. Module 6 (Venture) is surfaced last because it is the convexity anchor.

13 AUDIT, REPRODUCIBILITY, VERIFICATION

13.1 Append-Only Persistence

Score persistence is structurally append-only. The score_snapshots, portfolio_simulations, webhook_delivery_log, and audit_events tables block UPDATE and DELETE at the database trigger level. Service-role keys cannot bypass.

13.2 Cryptographic Hashing

Every snapshot carries a SHA-256 hash over the framework version, methodology version, sorted JSON of inputs and outputs, and PDF byte stream when applicable.

13.3 Bitcoin Anchor (OpenTimestamps)

Snapshots may optionally be anchored on Bitcoin's blockchain via OpenTimestamps. The anchor proves a given hash existed before a given block height — without trusting ACI.

13.4 Public Verify Endpoint

Any party can verify a score by snapshot ID without authentication via the publicly accessible score verification endpoint. The endpoint returns verification level, score and band, content hash, framework version, and Bitcoin anchor status. Technical integration details are published at aethoncredit.com/methodology/verify.

13.5 Determinism Test

A SHA-256 determinism test runs in continuous integration. For each module, the engine is invoked 100 times with identical inputs; all 100 SHA-256 hashes must be byte-identical.

14 VERSIONING AND MODEL GOVERNANCE

14.1 Active Framework Version

ACI Framework v1.0 — active since May 2026.

14.2 Change Process

Methodology changes require written CRO sign-off. Material changes trigger 30-day subscriber notification. Backtesting and validation are documented per change. The determinism guarantee is preserved across versions. Re-scoring is event-driven — historical scores are not silently re-computed.

15 METHODOLOGICAL LIMITATIONS

ACI Framework v1.0 is a transparent, deterministic methodology. Institutional users should be aware of the following structural limitations when interpreting outputs.

15.1 Public Evidence Constraints

ACI scores are computed exclusively from publicly available data. Private disclosures, undisclosed liabilities, internal communications, and non-public regulatory correspondence are structurally outside the evidence set. Entities with low public disclosure density will receive lower evidence confidence ratings and activate the Worst-Case Protocol for affected criteria. ACI does not represent that its evidence set is complete.

15.2 Historical Reconstruction Constraints

Calibration rests on reconstructed historical events using public evidence available at the time. Reconstruction involves judgment in evidence selection and bucket assignment. The dataset is not a formal statistical sample and does not support inference beyond the observed cases. Results should be interpreted as indicative of methodology performance on the observed cases, not as a guarantee of future performance. The calibration dataset will be expanded as additional failure events are documented and as live scoring data accumulates.

15.3 Outputs Are Not Forecasts

ACI outputs reflect structural risk at the time of computation. They are not predictions of default, failure, or price behaviour. Sudden fraud, undisclosed liabilities, or rapid structural deterioration may occur without prior signal in publicly available data. A LOW band indicator does not imply safety; a HIGH band indicator does not imply imminent failure. Outputs are one input among others in an independent analytical process.

15.4 Low-Disclosure Environments

Private entities, offshore-domiciled operators, and protocols without third-party audits present structurally lower evidence confidence. The Worst-Case Protocol substitutes conservative scores for low-confidence criteria, which depresses scores in low-disclosure environments. This is a design feature, not a deficiency: opacity is itself a risk signal. Users should treat a score driven predominantly by worst-case substitutions differently from a score driven predominantly by primary evidence.

15.5 Model Evolution and Known Gaps

ACI Framework v1.0 methodology gaps identified through the calibration process include: (i) absence of a non-recourse SPV exposure criterion in M4, which may understate financing risk isolated to special purpose vehicles; (ii) absence of a borrower disclosure quality criterion in M7 uncollateralised lending pools; (iii) absence of an oracle manipulation susceptibility criterion for DeFi derivatives protocols in M5 and M8; and (iv) the NULL criterion handling convention, which in v1.0 uses raw weighted sum rather than normalised-over-covered-weight for entities with undisclosed criteria. Each gap is documented in the methodology governance log and will be addressed in ACI Framework v2.0 subject to CRO approval and 30-day subscriber notification.

16 REFERENCE CALIBRATION

Illustrative worked example demonstrating end-to-end score computation for one Module M1B (Treasury Preferred) instrument under stated assumed inputs. Uses publicly observable data as of May 2026 and applies the full computational pipeline.

Illustrative — Not a Published Score

This worked example is provided as a methodology demonstration. Computed from stated assumed inputs under ACI Framework v1.0. Actual published scores will differ as evidence packs and market data evolve.

16.1 Worked Example — STRF (12-Month Duration)

Stated assumed inputs: Issuer: Strategy Inc. BTC holdings: 762,099 BTC. BTC price: $85,000. Senior debt: ~$8.21B. Aggregate preferred obligations: $2.10B.

```
SDACR = (762,099 × 85,000 − 8,210,000,000) ÷ 2,100,000,000  ≈  26.9×  →  Extreme tier  →  100
Income type: Fixed Contractual  →  100
Market Risk: round(80×0.50 + 100×0.30 + 75×0.20)  =  85
Raw = 100×0.30 + 100×0.25 + 85×0.20 + 100×0.17 + 100×0.08  =  97
Final = clamp( round( 97 × 1.10 ), 0, 100 )  =  100  →  LOW band  ·  Convexity: NEUTRAL
```

17 LEGAL DISCLAIMERS

17.1 Service Classification

ACI provides quantitative analytics and scenario computations to support independent analysis. It does not provide investment advice, portfolio management, or personal recommendations. ACI operates as a quantitative analytics and scenario computation platform providing model-generated risk metrics and scenario outputs derived from user-defined inputs and publicly available data.

17.2 Non-Reliance

ACI outputs are produced under a defined computational framework and not issued as credit ratings. Any decision is made independently by the recipient or in consultation with a licensed financial adviser. The recipient is a professional or institutional user capable of independently evaluating quantitative analytics outputs.

17.3 No Credit Rating Activity

ACI Risk Indicators and Provider Risk Profiles are computational outputs under a published framework and are not issued as credit ratings. They do not constitute "credit ratings" within the meaning of Regulation (EC) No 1060/2009 (EU CRA Regulation), the UK Credit Rating Agencies (Amendment etc.) (EU Exit) Regulations 2019, or any equivalent applicable law. ACI does not express opinions on the creditworthiness of any entity, instrument, or obligation.

17.4 Data and Methodology

ACI relies on publicly available and third-party data sources. Such data may be incomplete, delayed, or subject to revision. ACI applies its published methodology (ACI Framework v1.0, available at aethoncredit.com/methodology) consistently across all outputs. ACI does not guarantee the completeness, accuracy, or timeliness of any data or output.

17.5 No Advisory or Fiduciary Relationship

No fiduciary, advisory, or client relationship is created through access to or use of the ACI platform. ACI operates as an independent analytics provider.

User Acknowledgment Gate

I confirm that I am accessing Aethon Credit Intelligence as a professional or institutional user. I understand that ACI provides quantitative risk analytics and scenario computations designed to support independent analysis. I will apply ACI outputs as one input among others within my own analytical and governance processes, and I retain full responsibility for how they are used.

18 DOCUMENT CONTROL

Field	Value
Document title	ACI Risk Methodology
Document version	1.0 — Full Criterion-by-Criterion Edition
Framework version	ACI Framework v1.0
Issued	May 2026
Issuer	Aethon Credit Intelligence
Review cycle	Annual or material event trigger
Next scheduled review	May 2027
Approval	CRO sign-off required for any modification
Distribution	Public — published at aethoncredit.com/methodology
Source code reference	src/yield_engine/ — module-specific scoring.ts and bucket.score.map.ts
Disclaimer constants reference	src/lib/legal/disclaimer.constants.ts
Supersedes	N/A

ACI computes risk. You decide how to allocate capital.

This document is the authoritative public methodology disclosure for Aethon Credit Intelligence. Inconsistencies between this document and the running platform are reportable defects — contact help@aethoncredit.com