



USERS GUIDE & TRAINING

Aethon Credit Intelligence

Platform overview, module catalogue, and operating workflows

DOCUMENT

Users Guide + Training

FRAMEWORK

ACI v1.0

ISSUED

May 2026

CLASSIFICATION

Public Reference

QUANTITATIVE RISK ANALYTICS OUTPUT

What is Aethon Credit Intelligence?

ACI is a quantitative analytics and scenario computation platform. It computes Risk Indicators for digital asset yield-bearing positions across eight modules — including a stablecoin module (M1C) structured into two sub-modules — lets institutional users construct portfolio scenarios, runs deterministic stress tests, and produces auditable PDF reports.

Quantitative Scoring

Every provider scored 0–100 against the published bucket map. Deterministic, version-stamped, and hashed using SHA-256.

Scenario Computation

Eight modules, including a stablecoin module structured into two sub-modules. User constructs the scenario. ACI computes against ACI Framework v1.0.

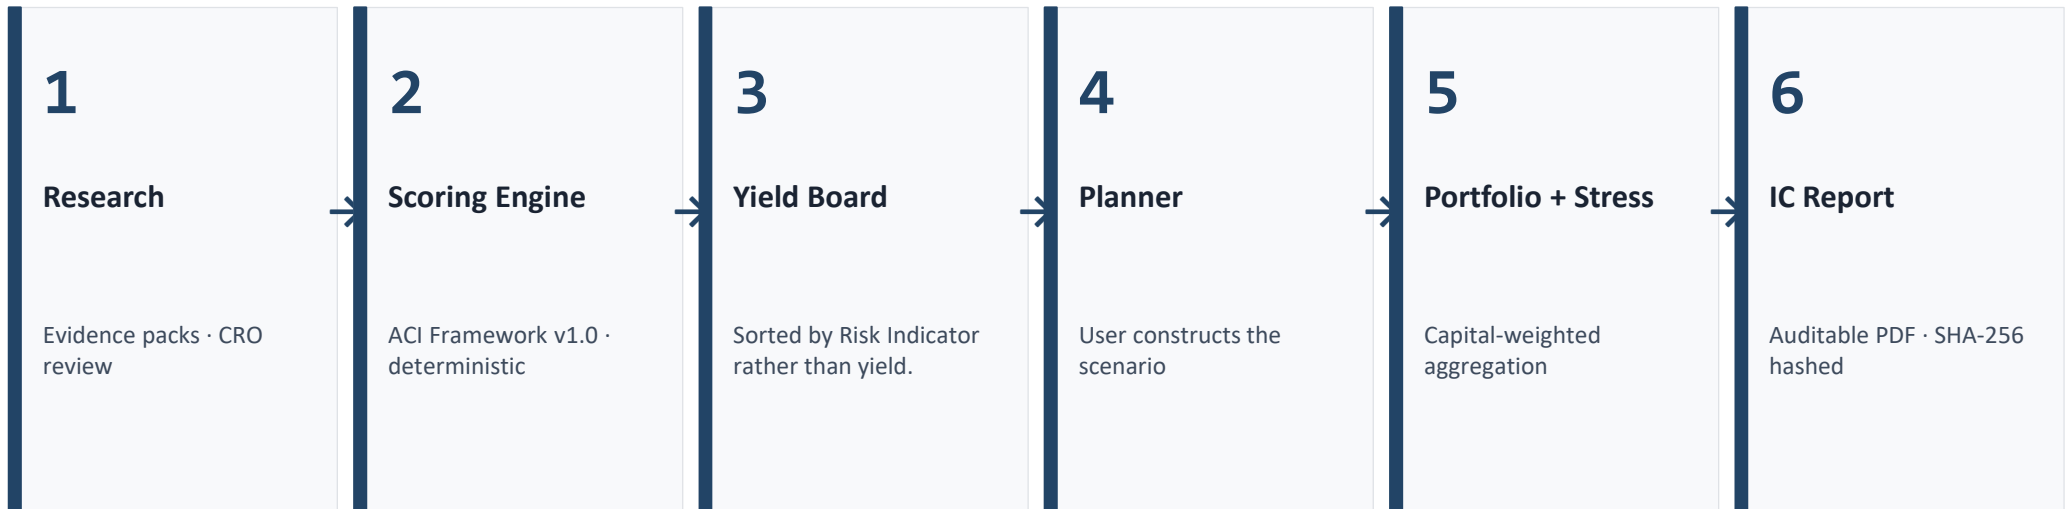
Auditable Outputs

Append-only score snapshots. Optional Bitcoin anchor via OpenTimestamps. Public verify endpoint.

ACI does not provide investment advice, portfolio management, execution services, or credit ratings. The user retains decision authority.

Platform at a Glance

The end-to-end flow from research and scoring to scenario, portfolio, stress, and report.



Every Risk Indicator is deterministic — identical inputs produce identical outputs across versions. Risk Indicators persist in append-only storage with cryptographic content hashes.

Service Classification

WHAT ACI IS

Aethon Credit Intelligence (ACI) operates as a quantitative analytics and scenario computation platform.

ACI provides model-generated risk metrics and scenario outputs derived from user-defined inputs and publicly available data.

System Boundaries

ACI does not provide investment advice, portfolio management, execution services, or credit ratings.

ACI does not express opinions on the creditworthiness of any entity, instrument, or financial obligation.

Source: ACI Legal Compass v2.0 · Block 3 — Service Classification (verbatim).

Three System Boundaries

LINE 1

No recommendation output

ACI outputs describe computed risk metrics for user-constructed scenarios. The platform never states which provider to choose, which allocation to use, or which action to take.

LINE 2

No creditworthiness opinion

ACI computes risk indicators derived from user-supplied evidence and publicly available data. ACI does not express a view on whether any provider will meet its financial obligations.

LINE 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.

Reference Scenarios are labelled "ACI Analytical Baseline (User Inputs Required)" and require explicit user confirmation of inputs before outputs are generated.

How the Scoring Engine Works

Each provider is evaluated against criteria defined in the published methodology for its module. Bucket values map to numeric scores via the deterministic Bucket Score Map.

- 1 Evidence pack assembled and CRO-reviewed
- 2 Bucket values mapped (e.g. 'tier_1', 'high', 'monthly')
- 3 Buckets resolve to 0–100 scores via the Bucket Score Map
- 4 Weighted average computed → raw score
- 5 Cascade penalty: three or more criteria scoring below 40 → minus 5 points
- 6 Hard caps applied (algo cap, depeg override, HV30 cap)
- 7 Duration multiplier applied to risk score only

Risk Bands and the ACI Risk Indicator

The ACI Risk Indicator is an integer score binned into four bands. Bands are computed deterministically. They are not credit ratings within the meaning of EU Credit Rating Agencies Regulation, the UK equivalent, or any applicable law.

LOW	80 - 100	Composite score in the upper quartile of the deterministic scale.
MEDIUM	60 - 79	Composite score in the second quartile.
ELEVATED	40 - 59	Composite score in the third quartile.
HIGH	0 - 39	Composite score in the lowest quartile, or hard-cap fired.

Hard caps may force a band reclassification regardless of composite score: algorithmic stablecoin → max 20; high depeg → forced HIGH; HV30 > 0.35 → market risk capped at 5.

Hard Caps and Cascade Penalty

RULE	TRIGGER	EFFECT
ALGO_STABLECOIN_CAP	Stablecoin peg_type = Algorithmic	Final score capped at 20
HIGH_DEPEG_OVERRIDE	Depeg > 1.5% over 90 days	Forced HIGH band regardless of score
HV30 cap (CRO 2026-04-30)	HV30 > 0.35	Market risk criterion capped at 5
CASCADE_PENALTY_APPLIED	Three or more criteria scoring below 40	Minus 5 to raw score before duration
VOLATILITY_REGIME_ACTIVE	HV30 / HV1Y regime multiple > 1.5x	Market risk shifts to HV30-dominant
ALL_CRITERIA_NOT_APPLICABLE	All criteria flagged not_applicable	Raw score → 0 (data integrity guard)

Hard caps are recorded in the score snapshot's hardCapsApplied[] array for full audit traceability. All rules are applied deterministically without manual adjustment.

Duration Multiplier

Longer commitment durations carry slightly higher risk weight. The duration multiplier adjusts the risk score only — never the projected income.

DURATION (MONTHS)	MULTIPLIER	USE CASE
≤ 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

CRITICAL RULE

Duration multiplier applies to risk score only — never to projected income.

Score Commentary — Deterministic and Auditable

Every score is accompanied by a deterministic commentary block explaining which criteria contributed to the result. All commentary is generated deterministically from inputs at computation time. Same inputs → same explanation.

EXAMPLE — STRF (TREASURY PREFERRED) • SCORE 74 MEDIUM

STRF scores 74 MEDIUM. The ceiling is held down by two factors: BTC Coverage Ratio (SDACR at 7.3× but senior debt reduces effective coverage, costing approximately 8 points) and Market Volatility (HV30 at 0.32 nearing the 0.35 hard cap threshold, costing approximately 6 points).

Ceiling factors

The 1–2 criteria dragging the score down most, with approximate point cost

Floor factors

Criteria providing the strongest support to the score

Band gap and duration

Distance to next band; pre- and post-multiplier scores when applicable

Fully deterministic.

Eight Modules (Stablecoin structured into two sub-modules)

PHASE 1 — YIELD-BEARING COUNTERPARTY COVERAGE

M1A

BTC Collateral Lending

Custody, LTV, jurisdiction, track record

M1B

Treasury Preferred Shares

BTC coverage ratio, HV1Y, convertibility

M1C-A

CeFi Stablecoins (M1C-A)

Reserves, audit, custody, withdrawal speed

M1C-B

DeFi Stablecoins (M1C-B)

Audit depth, immutability, governance, TVL

PHASE 2 — STRESS, PORTFOLIO, AND CONVEXITY

M4

Infrastructure

Mining + validator unit economics

M5

Market Neutral

Basis, funding, hybrid strategies

M6

Venture

Power-law return distributions

M7

Tokenised RWA

Credit tier, enforceability, custody

M8

Volatility

Options and structured products

M1A

BTC Collateral Lending

Custodial and non-custodial lending against BTC collateral. The user lends or pledges BTC; the platform issues credit or yield. ACI scores how the platform handles custody, top-up and liquidation, jurisdiction, and track record.

Representative providers: Ledn · Unchained · AMINA Bank · Sygnum · Strike · Figure

CRITERIA AND WEIGHTS

CRITERION	WEIGHT	WHAT IT MEASURES
Custody and rehypothecation	25%	Segregated vs pooled, disclosure of rehypothecation
Liquidation LTV	20%	Threshold at which collateral is force-sold
Top-up speed	15%	Hours from margin alert to required top-up
Jurisdiction tier	20%	Regulator and jurisdictional accountability
Maturity track record	20%	Years operating, exploit history, prior workouts

M1B

Treasury Preferred Shares

BTC-treasury preferred share instruments. Companies holding BTC reserves issue preferred shares paying fixed dividends. ACI scores BTC-coverage ratio (SDACR), volatility, convertibility, and instrument income mechanics.

Representative providers: STRF · STRK · STRC · STRD · SATA · STRE

CRITERIA AND WEIGHTS

CRITERION	WEIGHT	WHAT IT MEASURES
BTC coverage ratio (SDACR)	30%	Holdings × price minus senior debt, divided by preferred obligations
Income mechanism	20%	Fixed contractual vs variable / discretionary dividend
Market risk composite	20%	HV1Y, price-to-par, 30-day volume — three components
Convertibility	15%	Forced equity conversion risk under market stress
Issuer maturity	15%	Institutional vs early-stage issuer profile

M1C-A

CeFi Stablecoins

Centralised stablecoin yield platforms. The user deposits stablecoins; the platform earns yield via lending, market-making, or treasury operations. ACI scores reserve disclosure, audit depth, custody, and jurisdictional accountability.

Representative providers: Coinbase · Kraken Earn · Nexo · OKX Earn · Matrixport · YouHodler

CRITERIA AND WEIGHTS

CRITERION	WEIGHT	WHAT IT MEASURES
Platform reserve tier	25%	Segregated platform reserves vs commingled
Yield source + disclosure	25%	Contractual vs discretionary; disclosure quality
Custody disclosure	20%	Segregated and audited vs unaudited self-custody
Withdrawal speed	15%	Instant vs same-day vs delayed
Jurisdiction + peg type	15%	Regulator tier and underlying peg mechanism

M1C-B

DeFi Stablecoins

On-chain stablecoin yield protocols. Smart-contract platforms paying yield from over-collateralised lending, derivatives basis, or RWA exposure. ACI scores audit depth, immutability, governance, transparency, and depeg trajectory.

Representative providers: Aave v3 · Compound v3 · Spark Protocol · Morpho Blue · Ethena USDe

CRITERIA AND WEIGHTS

CRITERION	WEIGHT	WHAT IT MEASURES
Audit depth tier	25%	Tier-1 multi-audits vs single audit vs no audit
Battle-test tier	20%	Years operating without exploit
Immutability	20%	Timelock duration and governance controls
Yield transparency	15%	Fully on-chain documented vs opaque
TVL + withdrawal	20%	Liquidity depth and withdrawal speed

M4

Infrastructure

Bitcoin mining and validator operators. Yield from block rewards, transaction fees, or staking. ACI scores unit economics under BTC price shocks, energy cost structure, uptime, and jurisdictional exposure.

Representative providers: Marathon · Riot · Cipher Mining · Hut 8 · Core Scientific · Compass

CRITERIA AND WEIGHTS

CRITERION	WEIGHT	WHAT IT MEASURES
Breakeven BTC price	20%	Distance to monthly cashflow zero
Energy cost / kWh	15%	Per-kWh cost benchmarked to industry distribution
Margin stability	15%	Variability of net margin across recent quarters
Uptime + maintenance	20%	Realised uptime and maintenance quality tier
Legal + custody	15%	Asset custody tier and legal structure
Jurisdiction + audit	15%	Tier-1 transparent vs offshore, audit cadence

M5

Market Neutral

Basis, funding-rate, and hybrid market-neutral strategies. Yield from harvesting persistent funding-rate or basis-trade premiums. ACI scores strategy robustness, risk management, execution quality, liquidity access, and transparency.

Representative providers: Wintermute · GSR · Cumberland · Amber Group · B2C2

CRITERIA AND WEIGHTS

CRITERION	WEIGHT	WHAT IT MEASURES
Strategy + Sharpe + AUM	30%	Strategy type tier, Sharpe ratio tier, AUM disclosure
Risk management	25%	Leverage, counterparty, collateral segregation
Execution quality	20%	Funding-rate access, venue tier, pricing infrastructure
Liquidity access	15%	Exchange diversity, position visibility
Transparency + audit	10%	Track record, drawdown history, disclosure, audit

M6

Venture

Digital asset venture funds. Power-law return distributions; long holding periods; concentrated outcomes. ACI surfaces IRR bands, vintage concentration, time horizon, and the optionality / convexity profile of the fund subtype.

Representative providers: a16z crypto · Polychain · Pantera · Paradigm · Dragonfly

CRITERIA AND WEIGHTS

CRITERION	WEIGHT	WHAT IT MEASURES
Fund subtype	25%	Multi-stage VC vs early-stage VC vs liquid fund
IRR bands (low/mid/high)	25%	Distribution-based outcome estimates
Holding period	15%	Time-horizon mismatch with investor capital
Vintage concentration	15%	Concentration in single vintage years
GP track record	20%	Historical fund performance and DPI

M7

Tokenised RWA

On-chain real-world asset platforms — sovereign bills, corporate credit, private credit, structured products. ACI scores credit tier, enforceability, custody, secondary-market depth, and reporting cadence.

Representative providers: Securitize · Backed Finance · Ondo Finance · Maple Finance · Centrifuge

CRITERIA AND WEIGHTS

CRITERION	WEIGHT	WHAT IT MEASURES
Credit tier (IG / SIG / PC)	30%	Defaults: pc.default=0.035, pc.recovery=0.60
Legal structure	25%	Enforceability, investor protection, jurisdiction
Custody + bankruptcy	20%	Qualified custodian, bankruptcy-remote SPV
Liquidity terms	15%	Daily / periodic / locked redemption
Reporting + audit	10%	Real-time vs periodic; audited tier

M8

Volatility

Options, structured products, and DeFi option vaults. Convexity-aware scoring across exchange-venue, institutional-OTC, and on-chain-DOV segments. ACI scores strategy design, counterparty, execution, track record, and transparency.

Representative providers: Deribit · Wintermute (OTC) · Ribbon Finance · Friktion · Oryn

CRITERIA AND WEIGHTS

CRITERION	WEIGHT	WHAT IT MEASURES
Strategy design	25%	Defined-risk vs naked-directional
Counterparty quality	25%	Tier-one regulated vs unrated offshore
Execution infrastructure	20%	Multi-venue smart routing vs manual OTC
Track record	15%	Years operating without major drawdown
Transparency	15%	Monthly full positions vs opaque

Getting Started — Your First Scenario in Five Steps

1**Browse the Yield Board**

All scored providers ranked by ACI Risk Indicator (never by APY). Filter by module.

2**Choose a planner module**

Select one of M1A / M1B / M1C (CeFi: M1C-A · DeFi: M1C-B) / M4 / M5 / M6 / M7 / M8.

3**Construct your scenario**

Use Reference (auto-selects from the highest-scoring providers) or My Scenario (full provider control).

4**Run stress scenarios**

Apply the module's stress engine. Add custom slider stresses if your IC requires it.

5**Save and export**

Save scenarios. Aggregate into a portfolio. Generate the IC PDF report.

Two Ways to Construct a Scenario

REFERENCE SCENARIO

Analytical Baseline (User Inputs Required)

- ACI auto-selects providers ranked by Risk Indicator
- Default reference weighting: 40 / 30 / 30 (1C uses 70 / 15 / 15)
- Only providers in LOW or MEDIUM band eligible by default
- User must confirm or modify inputs before outputs render
- Never presented as a default strategy

MY SCENARIO

Custom Construction — Full Provider Control

- User selects any provider regardless of band
- Weights set by the user; validated in real time
- ELEVATED-band providers capped at 30% allocation
- HIGH-band providers gated by full-screen acknowledgment
- Used to evaluate sensitivity to alternative provider selections.

Tip: open Reference to see the baseline, then switch to My Scenario to model deviations.

Inside the Planner

Each module's planner page exposes the same four-stage flow. Inputs are user-supplied; outputs are deterministic.

01

Inputs

Capital, duration, allocation. Module-specific inputs (SDACR ticker, peg type, strategy type) surface conditionally.

02

Compute

Engine applies ACI Framework v1.0. Bucket lookup → weighted score → cascade penalty → hard caps → duration.

03

Output

ACI Risk Indicator, scenario score, criterion breakdown, score commentary, optional convexity classifier.

04

Persist

Save to localStorage (anonymous) or planner_sessions DB row. Append-only at the database trigger level.

Reading the Output Stack

Every scenario produces a structured output stack. Each element is computed deterministically from user inputs.

ACI Risk Indicator

Integer 0–100, binned into LOW / MEDIUM / ELEVATED / HIGH. Computed against ACI Framework v1.0 bucket map. Not a credit rating.

Scenario score and criterion breakdown

Composite of all criterion scores weighted per module. Cascade penalty and hard caps applied before duration. Per-criterion breakdown surfaces inputs that drove the result.

Score commentary

Deterministic narrative naming the 1–2 ceiling factors and 1–2 floor factors with approximate point cost. No model inference.

Hard caps applied[]

List of any cap or modifier that fired (ALGO_STABLECOIN_CAP, HIGH_DEPEG_OVERRIDE, CASCADE_PENALTY_APPLIED, etc.).

Capital-Weighted Portfolio Aggregator

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.

$$\text{Portfolio Score} = \frac{\sum (\text{score}_i \times \text{capital}_i \times \text{df}(\text{duration}_i))}{\sum (\text{capital}_i \times \text{df}(\text{duration}_i))}$$

DURATION FACTOR $\text{df}(\text{duration})$

MONTHS	FACTOR
≤ 3	1.000
≤ 6	1.050
≤ 12	1.100
≤ 24	1.175
> 24	1.250

INVARIANTS

- Activation gate: ≥ 2 sessions across ≥ 2 distinct modules
- Duration multiplier applies to risk score only, never income
- Final score clamped to $[0, 100]$, rounded to integer
- Scores read from locked snapshots — never re-computed
- Canonical module order: 1A \rightarrow 1B \rightarrow 1C \rightarrow 4 \rightarrow 5 \rightarrow 7 \rightarrow 8 \rightarrow 6 (M1C encompasses sub-modules 1C-A and 1C-B)

Cross-Module Counterparty Concentration

When a single counterparty appears across multiple modules — for example, the same custodian in M1A and M1C — combined exposure is computed and surfaced.

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 %

The platform also detects DeFi-issuer correlation (e.g. multiple USDC-backed protocols) and Module 8 short-volatility segment concentration.

Stress Engines for M4 – M8

Each Phase 2 module ships a stress engine that projects outputs against pre-defined adverse scenarios, plus user-defined custom-slider scenarios.

M4 Energy shock · difficulty drift · combined adverse · halving compression

M5 Rate compression · correlation shock · extreme drawdown · liquidity crunch

M6 Vintage concentration · macro contraction · correlation shock · early-redemption stress

M7 Spread widening · sovereign stress · liquidity trigger · combined adverse

M8 Spot drawdown · time decay · vol regime shift · combined adverse

Stress outputs are scenario-specific computations — they are not forecasts, predictions, or recommendations.

IC Reports — Auditable PDF Output

Every saved scenario or portfolio can produce a paginated IC PDF report. Reports inherit the platform design system and are constructed deterministically from the locked snapshot.

Page 1 — Cover

Report ID, subject, framework version

Page 2 — Evidence Summary

Risk band, composite score, confidence, drivers, constraints

Page 3 — Scope

What this report covers and explicit out-of-scope statements

Pages 4-6 — Analysis

Dimension breakdown, stress analysis, concentration, history

Pages 7-8 — Constraints + Comparison

Policy flags and peer / category context

Pages 9-10 — Appendix + Legal

Data sources, version tags, verbatim disclaimer

Append-Only Audit + Bitcoin Anchor

Score persistence is structurally append-only. Every snapshot stamps the active framework version, computes a SHA-256 content hash, and may optionally anchor that hash on Bitcoin via OpenTimestamps.

Append-Only Database

score_snapshots blocks UPDATE and DELETE at the database trigger level. Same enforcement applies to portfolio_simulations, webhook_delivery_log, and audit_events. service_role keys cannot bypass these controls.

Deterministic Engine

Same inputs → same outputs across processes. Verified by SHA-256 determinism tests covering all eight modules (including both M1C sub-modules). Methodology updates are version-controlled; historical scores remain reproducible.

Cryptographic Hash

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.

Bitcoin Anchor

OpenTimestamps proof generated and committed to Bitcoin's blockchain. Allows external parties to verify a score was computed before a given block height — without trusting ACI.

Score Verification — Public Endpoint

Any party — institutional user, external auditor, regulator — can verify a score by snapshot ID without authentication.

GET /api/verify/{scoreId}

RESPONSE FIELDS

verified	Boolean — score located in append-only store
verificationLevel	VERIFIED PENDING LOCAL_ONLY MISMATCH
score · band · hash	Integer 0–100, band label, SHA-256 of snapshot content
scoredAt · framework	ISO timestamp and framework version stamped at compute time
anchorStatus · hardRules	Bitcoin anchor state plus list of hard caps that fired

Methodology Versioning

The published computational framework is version-controlled. Every score is stamped with the active framework version at compute time. Historical scores remain reproducible under their original version.

CURRENT VERSION

ACI Framework v1.0

Active since 22 April 2026.

VERSION RULES

- CRO sign-off required
- 30-day subscriber notification for material changes
- Backtesting + validation documented
- Determinism guarantee preserved

PUBLIC METHODOLOGY DISCLOSURE

Full methodology — every weight, every threshold, every bucket map, every hard cap — is published at aethoncredit.com/methodology. The platform's outputs are reproducible from disclosure alone.

Source: ACI Legal Compass v2.0 · Block 7 — Data and Methodology.

User Acknowledgment Gate

First-time access to the ACI platform requires a one-time acknowledgment. The text below is presented as a click-through gate. The acknowledgment is logged with user ID, timestamp, and platform version.

BLOCK 13 – VERBATIM

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.

Source: ACI Legal Compass v2.0 · Block 13 (verbatim). Exported as DISCLAIMER_ACKNOWLEDGMENT_GATE in src/lib/legal/disclaimer.constants.ts.

Quick Reference — Key Terms

ACI Risk Indicator

Integer 0–100, computed deterministically against the published bucket map. Not a credit rating.

Bucket Score Map

Lookup table converting categorical bucket values to numeric scores 0–100. Versioned.

Duration multiplier

Scaling factor 1.0 to 1.25 by duration band. Applied to risk score only — never to projected income.

Reference Scenario

Analytical Baseline (User Inputs Required). Auto-selected providers; user must confirm before output.

ACI Risk Band

LOW (80–100) · MEDIUM (60–79) · ELEVATED (40–59) · HIGH (0–39). Hard caps may force reclassification.

Cascade penalty

Minus 5 to raw score when three or more criteria score below 40. Applied before duration multiplier.

Hard cap

Engine-level override that constrains the score (e.g. ALGO_STABLECOIN_CAP = 20).

Append-only

DB enforcement — UPDATE and DELETE blocked at the trigger level on score_snapshots and related tables.

AETHON

CREDIT INTELLIGENCE

Quantitative analytics. Deterministic computation. Auditable outputs.

METHODOLOGY

aethoncredit.com/methodology

VERIFY A SCORE

aethoncredit.com/verify/{scoreId}

CONTACT

help@aethoncredit.com

Aethon Credit Intelligence · ACI Framework v1.0 · May 2026 · 34 / 34

QUANTITATIVE RISK ANALYTICS OUTPUT