

TVAC Methodology & Calibration v1

Technical note – Version 1.2 (4 March 2026)

Purpose: A non-marketing technical description of TVAC's fixed evaluation framework, scoring logic, verdict mapping, calibration approach, and key limitations.

Contents

- 1. Purpose and scope
- 2. System overview (process and architecture)
- 3. Conceptual model and Added Value formula
- 4. Inputs, validation, and assumptions
- 5. The five factors (0–10 scale)
- 6. Scoring workflow, normalization, and report generation
- 7. Verdict bands and thresholds
- 8. Stakeholder and initiator economics (qualified / constrained)
- 9. Calibration, regression checks, and reference cases
- 10. Limitations and appropriate use

1. Purpose and scope

The Tokenization Value-Add Calculator (TVAC) is an evaluation engine for tokenization cases. It is designed to answer one practical question: **Does this specific tokenization design create net added value compared to a conventional structure?** TVAC is an informational decision-support tool. It does not provide legal, tax, financial, or investment advice.

TVAC is intentionally conservative. When a case sits on a boundary, the more cautious interpretation is typically chosen.

2. System overview (process and architecture)

TVAC is a method-governed pipeline: the AI model is constrained by a fixed framework and strict output requirements.

- You describe the case in free text and answer structured questions.
- The app validates inputs and packages them into a single case payload.
- On 'Generate report', the payload is sent securely to the TVAC backend (API).
- The backend applies TVAC's fixed methodology: Added Value formula, criteria, and report structure.
- A tightly structured request (strict format and evaluation rules) is sent via API to an AI model.
- The model returns a structured analysis (verdict, rationale, factor breakdown, risks, recommendations).
- TVAC parses, consolidates, and sanity-checks the output against required sections and consistency rules.
- Results are normalized into a report-ready structure (chapters, tables, visuals).
- The API returns the full report content/data to the browser; the UI renders the final report and supports print-to-PDF.

Important: TVAC does not publish raw model output. The model is constrained by TVAC's methodology and required structure; results are checked and normalized before rendering.

3. Conceptual model and Added Value formula

Each case is evaluated across five core factors on a 0–10 scale:

- **New Opportunities (NO)**
- **Cost Savings (CS)**
- **Risk Reduction (RR)**
- **Tokenization Costs (TC)**
- **New Risks (NR)**

TVAC combines the factor scores into a single Added Value number using:

$$\text{Added Value} = (\text{NO} + \text{CS} + \text{RR}) - (\text{TC} + \text{NR})$$

In shorthand: **Plus** = NO + CS + RR; **Minus** = TC + NR; **Added Value** = Plus - Minus.

4. Inputs, validation, and assumptions

TVAC uses two categories of inputs:

- **Free-text case description** (recommended 150–600 words; max ~900). Describe what is being tokenized, who the investors are, the regulatory path, custody/settlement, liquidity/venue design, objectives, and constraints.
- **Structured fields** that anchor the narrative (issuance jurisdiction, sale jurisdictions, investor types, instrument type, custody model, liquidity cadence and venue type, settlement model, disclosure path, transfer restrictions, optional tax notes).

TVAC does not scrape the web or query hidden external databases at runtime. The assessment is based on the information provided by the user and TVAC's fixed rubrics.

The report explicitly lists key assumptions and missing inputs where relevant. Thin or contradictory inputs should not be treated the same as a carefully specified case.

5. The five factors (0–10 scale)

Each factor is scored from 0 to 10. For upside factors (NO, CS, RR), higher scores are better. For friction factors (TC, NR), higher scores are worse.

5.1 New Opportunities (NO)

Case-specific upside that is only possible, or meaningfully strengthened, by tokenization compared to a conventional setup.

- Access to new investor segments (within a credible compliance path).
- More realistic secondary liquidity or new market structure capabilities.
- Programmability that removes real frictions (not cosmetic wrappers).

5.2 Cost Savings (CS)

Structural lifecycle cost reduction vs a well-run traditional baseline (not vs a dysfunctional baseline).

- Reduced reconciliation and break management.
- Automation of lifecycle events and reporting.
- Fewer manual steps and duplicated ledgers across intermediaries (where realistically removed).

5.3 Risk Reduction (RR)

Reduction in operational, counterparty, governance or transparency risk compared with baseline operations.

- Stronger auditability and tamper-evident records.
- More robust delivery-versus-payment (DvP) mechanics and fewer settlement fails.
- Clearer control points for transfer restrictions and eligibility.

5.4 Tokenization Costs (TC)

Incremental costs introduced by tokenization beyond a conventional structure. Higher is worse.

- Legal and structuring work, plus ongoing compliance burden.
- Platform/custody/venue integration and operational change management.
- Smart-contract development/audit where applicable.

5.5 New Risks (NR)

Incremental risks created by tokenization vs baseline. Higher is worse.

- Technology / smart-contract / integration risk.
- Custody, key management, and provider concentration risk.
- Regulatory classification and market-structure uncertainty.

6. Scoring workflow, normalization, and report generation

At a high level, TVAC evaluates a case in this sequence:

- Interpret the case based on free text plus structured fields.
- Identify mechanisms relevant to each factor (drivers of upside and sources of friction).
- Assign factor scores using fixed rubrics; compute Plus, Minus, and Added Value.
- Apply deterministic verdict mapping and stop criteria.
- Generate a structured report and run normalization/sanity checks before returning output to the UI.

Normalization and consistency checks. TVAC runs post-processing checks that ensure internal consistency between numerical scores and narrative text (e.g., references to TC/NR/Added Value in recommendations).

Non-viable framing. For Non-viable outcomes, TVAC explicitly frames improvement guidance as redesign requirements (or reasons to stop), not as 'tuning levers'.

7. Verdict bands and thresholds

TVAC derives a verdict using deterministic thresholds on Added Value, combined with stop criteria for structurally impossible or contradictory designs.

Indicative thresholds (current calibration):

- **Viable:** Added Value ≥ 4
- **Borderline:** Added Value ≥ 2 (but < 4)
- **Conditional Go:** Added Value ≥ -1 (but < 2), or cases where the concept is directionally viable but depends on specific execution conditions.
- **Non-viable:** Added Value ≤ -2 , or presence of hard blockers/contradictions.

These thresholds are not tuned per user. They may evolve slightly between versions, but remain stable within a release.

8. Stakeholder and initiator economics (qualified / constrained)

TVAC includes an optional stakeholder added-value view (Asset Owner/Issuer, Investors, Infrastructure providers). These stakeholder net scores are heuristic and not time- or scale-discounted; they are intended as distribution and incentive-alignment insight, not as a definitive ROI model.

TVAC also emits an initiator economics qualifier to improve nuance for long-horizon institutional cases:

- **ok:** initiator net is non-negative (heuristic).

- **qualified:** initiator net is negative but plausibly improves over time and/or at scale, especially in post-trade modernization cases. Treat as a caveat that must be addressed (value capture, milestones, decommissioning legacy processes).
- **constrained:** initiator net is strongly negative without strong long-horizon/scale signals; near-term value capture must be made explicit.

9. Calibration, regression checks, and reference cases

TVAC is calibrated against a fixed library of synthetic reference cases (representative design patterns) and sanity-checked against a small set of historic real-world structures. The objective is consistency and decision usefulness, not perfect numerical accuracy on any single case.

Regression checks. Each release is checked to avoid drift in verdict mapping and to ensure that known reference cases continue to land in the expected verdict bands.

10. Limitations and appropriate use

- TVAC is informational decision-support only. It is not legal, tax, financial, or investment advice.
- TVAC does not assume responsibility for outcomes. Users must validate jurisdiction-specific matters with qualified advisers.
- Scores are calibrated heuristics; they are not an accounting model and do not replace quantitative project economics.
- TVAC is not a marketplace, not an issuance platform, and not an underwriting or approval function.
- Where tokenization is legally or operationally impossible under the stated constraints, TVAC will mark the case Non-viable regardless of narrative ambition.