THE FRACTIONAL OPERATING STANDARD

FOS · Version 1.0

Raine Townsend · Designer of The Fractional Operating Standard · Founder, Operations Protocol · 2026

Fractional engagements don't end when the work gets bad. They end when authority erodes — quietly, structurally, and measurably — months before anyone says the word "transition." The Fractional Operating Standard exists to make that erosion visible while it can still be reversed, and to run the practice that catches it as a system instead of a feeling.

What this Standard is — and is not

The Fractional Operating Standard is an operating doctrine for solo fractional executives — COOs, CFOs, CROs, CTOs, CMOs running multiple concurrent engagements under retainer. It defines what a fractional practice must record, what it must compute from those records, and what decisions those computations must trigger.

It is not an academic study, and it does not pretend to be one. It is codified operating doctrine: the failure patterns are the ones practitioners recognize on sight, and every metric in this document is implemented, end to end, in a working reference system. Nothing here is aspirational. If the Standard says a number is computable from the inputs, it is computed from the inputs.

Nor was it written from inside a single fractional practice. It was engineered from outside the seat — deliberately. Practitioners meet this failure pattern one career at a time; an instrument requires the pattern isolated from any one career, weighted by structure rather than by memory, and proven by implementation before publication. That is the vantage this document is written from.

It is also not a course. The Standard assumes you already know how to do the work of your function. What it systematizes is the part almost nobody systematizes: the engagement around the work — scope, authority, evidence, economics, and the renewal decision.

Three commitments run through everything that follows:

Record decisions, not vibes. Every input in the Standard is something that happened — a score you assigned with written evidence, an indicator you marked with a date, a dollar figure you can defend.

Compute health; don't intuit it. Composite numbers exist so that drift is visible before it is dramatic. A practitioner's gut catches erosion at month five. The instruments catch it at week three.

Let thresholds force conversations. The Standard's tiers and windows exist to convert an awkward, avoidable conversation into a scheduled, structural one.

A note on notation: every metric in the Standard carries an index code — M1 through M18 — in its reference registry. Prose in this document uses the metrics' names; the codes appear in formulas and tables, where precision earns them.

---

1 · The failure model: authority erodes before revenue does

Ask a fractional executive why an engagement ended and you will usually hear a proximate cause — budget, reorganization, "they brought it in-house." Ask what the six months before that looked like and a different story emerges, and it is almost always the same story.

The engagement didn't die at the renewal conversation. It died in increments: a decision made inside the FE's scope without the FE in the room. Operational email routed around them. A reversal, delivered casually, of a call they'd made three weeks earlier. A meeting about their function that they learned about afterward. None of these is a crisis. Each is deniable. Together they form a pattern with a direction, and the direction is out.

The Standard names this pattern **authority erosion**, and rests on one operating thesis:

> ***Authority erosion precedes revenue loss, and erosion is measurable before it is sayable.***

Measurable — because each erosion event is a discrete, observable fact. Before it is sayable — because no single event justifies raising the alarm, which is exactly why practitioners raise it late or never.

The authority posture ladder

Engagement authority is not binary. The Standard describes it as a four-position ladder, and requires the position to be recorded, with dates, as it changes:

1. **Observer** — you attend, you read, you are briefed. No decision routes through you.

2. **Advisor** — your judgment is sought before decisions; the decisions remain elsewhere.
3. **Operator** — designated decisions are yours. You hold the pen inside a documented scope.
4. **Leader** — you own outcomes, direct people, and are accountable for the function's results.

Two rules attach to the ladder. First: **integration has a clock.** An engagement still sitting at Observer past day fourteen is not ramping — it is stalling, and the Standard treats it as an early governance signal, not a scheduling artifact. Second: **direction matters more than position.** An Operator sliding toward Advisor is a louder alarm than an Advisor who was hired to advise.

Named failure modes

The Standard tracks a small set of named, practice-level failure modes because naming them makes them checkable:

- **Single-Client Dependency** — one engagement exceeds 40% of practice revenue. The practice has an employer, not a portfolio, and every negotiation with that client happens under duress.
- **Scope Drift** — work accumulates outside the documented scope faster than change orders document it. Drift is unpaid work today and a renewal-time dispute later.
- **Silent Erosion** — the pattern of this section, running unmeasured.

Everything in the sections that follow exists to make these three impossible to miss.

---

2 · The canonical inputs: what a practice must record

The Standard defines a fixed registry of **eighteen canonical inputs** — the complete set of facts a fractional practice records, from which every metric in this document is derived. The registry is deliberately closed and deliberately small: if a number matters, it is computed from these; if it cannot be computed from these, the Standard does not claim it.

The inputs divide into two scopes.

Practice-level inputs — recorded once, about you:

- **Hour ceiling** — the total monthly hours the practice can sell. Not a goal; a constraint.
- **Administrative reserve** — hours held back for running the practice itself. Capacity math that ignores admin is fiction.
- **Expense ratio** — the percentage of gross that operating the practice consumes.

- **Target client count** — the portfolio size the practice is built for.

Engagement-level inputs — recorded per client, at the moments they become true:

- **Qualification decision** — proceed, proceed-with-conditions, or decline, recorded before the SOW exists.
- **Engagement tier** — the contracted intensity of the engagement (Section 7 defines the three tiers and their load multipliers).
- **Retainer** — the monthly number, exactly as contracted.
- **Committed hours** — the monthly hours the retainer buys.
- **Start and end dates** — because every engagement has an end date, and the renewal clock (Section 5) runs from it.
- **Audit scores** — the six-dimension health assessment of Section 3, scored at baseline, month three, and month six, each score carrying written evidence.
- **Scope change register** — every approved change order, dated. The trailing count is itself a governance signal.
- **Governance indicator states** — the ten-indicator register of Section 4, the Standard's core instrument.
- **Value ledger entries** — every quantified item of delivered value, with its implementation friction, per Section 6.
- **Renewal records** — every renewal decision and its outcome, in the closed vocabulary of Section 5.
- **Stakeholder map** — who holds influence, who holds disposition, and who your champion actually is.
- **Authority posture** — the ladder position of Section 1, with transition dates.

Two registry conventions are part of the doctrine. Inputs are **numbered and versioned** in the reference registry: one slot is currently reserved for future assignment, and one has been formally retired and superseded (the legacy renewal-outcome field, replaced by the renewal records above). A registry that can retire an input without renumbering the rest is a registry that can survive contact with reality.

The test of the registry is severe and intentional: **if a fact about an engagement is going to matter at renewal, it must live in a canonical input the day it happens — not in your memory, and not in your inbox.**

3 · Engagement health: the M1 composite

Every engagement in a Standard-run practice carries one number that answers "how is this engagement, structurally?" That number is the **Health Composite (M1)**, scored 0–100, and it is an audit result — not a mood.

The six dimensions

The composite is computed from six dimensions, each scored 0–10 against written anchors, each score accompanied by recorded evidence. The weights are fixed by the Standard:

Dimension	Weight	What it measures
Scope & Authority Clarity	**× 1.5**	Is what you own documented, agreed, and holding?
Stakeholder Alignment	× 1.4	Do the people who matter agree on why you're there?
Operational Complexity	× 1.2	Is the environment tractable, or does chaos tax every hour?
Decision Velocity	× 1.1	Do decisions you need actually happen, at speed?
Financial Visibility	× 1.0	Can you see the numbers your function requires?
Communication Cadence	× 1.0	Does information flow on a rhythm, or on accident?

M1 = Σ(dimension score × weight) ÷ 7.2 ÷ 10 × 100 — a weighted composite normalized to 0–100.

Why Scope & Authority carries 1.5×

The heaviest weight in the model is an argument, and it is the Standard's central one. Stakeholders can be re-aligned. Complexity can be managed down. Cadence can be rebuilt in a month. But an engagement where scope and authority have gone soft is an engagement where every other dimension is now negotiable — where "what you own" is decided ad hoc, meeting by meeting, by whoever is in the room. Scope & Authority Clarity is weighted 1.5× because it is the dimension the other five depend on, and because its decay is the earliest reliable predictor of the failure pattern in Section 1.

Cadence and evidence

The audit runs at **baseline, month three, and month six** — often enough to produce a trend, spaced enough that movement means something. Two disciplines make the number trustworthy:

Anchored scoring. Each dimension defines what a 2, a 5, and an 8 look like. You are matching observed conditions to described states, not expressing satisfaction.

Evidence or it didn't happen. Every score carries written observations. A score you can't evidence is a score you'll quietly inflate — and the audit exists precisely for the quarters when you'd rather not know.

The trend is the product. A 71 means little alone; 84 → 76 → 71 across three audits, with the Scope & Authority row driving the slide, is Section 1 happening on schedule — and Section 4 exists to catch it between audits.

4 · The Governance Shield: ten indicators, one score, four tiers

The audit of Section 3 runs at fixed marks. Erosion runs weekly. The **Governance Shield** is the continuous instrument between audits: ten binary indicators, each an observable event, each weighted by how directly it strikes at authority.

The ten indicators

Structural Authority — weight ×3. The events that are the erosion:

1. **Decisions Bypassed** — a decision inside your scope, made without your input. You were hired to provide judgment; decisions made without it defeat the engagement's purpose.
2. **Communication Routed Around You** — operational traffic directed to your reports or scope counterparts without you copied. You cannot operate on information you don't have.
3. **Decisions Reversed** — a call of yours reversed within thirty days, without consultation. Public reversals damage standing; standing is the asset.
4. **Parallel Solution Evaluated** — leadership evaluating an alternative to your approach — another consultant, a full-time hire, an internal option — without telling you. Sponsor confidence has already moved.

Visibility & Contract — weight ×2. The conditions that let erosion happen unrecorded:

5. **Scope Undocumented** — work performed outside the SOW without a written change order. Creep without an audit trail is a renewal dispute on layaway.
6. **Meeting Exclusions** — key operational meetings held without you, especially when your function is on the agenda.
7. **Budget Visibility Restricted** — you are asked to manage outcomes without seeing the inputs: a COO without OpEx access, a CFO without the revenue forecast.
8. **Sponsor / Champion Changed** — the executive who hired you has departed, been reassigned, or been diminished. The conditions your SOW was signed under may no longer exist.
9. **SOW Disputed** — scope, deliverables, or terms contested outside the change-order process. The contractual floor is moving.

Friction — weight ×1. The canary:

10. **Sponsor Response Past 72 Hours** — a written request requiring decision, unanswered past seventy-two hours, within the trailing fourteen days. Decision velocity is broken at the top.

The M17 score

The register produces the **Governance Score (M17)** — the early-warning number:

M17 = (1 − weighted "yes" ÷ 23) × 100

Twenty-three is the register's total weight — four ×3 indicators, five ×2, one ×1. All clear scores 100; every active indicator pulls the score down in proportion to how structural it is. Reading bands: **80 and above, healthy · 60–79, caution · 45–59, at risk · below 45, severe.**

The weighting is itself doctrine: a missed reply (×1) is friction; a restricted budget view (×2) is a condition; a bypassed decision (×3) is the disease. A single structural event outweighs the canary threefold because that is how it works in the field.

The four governance tiers

The Governance Score is a reading. The **tiers** are the required response, evaluated from the top down:

Tier 4 — Engagement at Risk. Triggered by a Governance Score below 40, two or more Structural Authority indicators active, a sponsor change compounding any structural indicator — or a single confirmed violation of your documented sole or lead decision authority. That last trigger is absolute by design: a breach of agreed structure is not a data point to weigh; it is the structure failing. Tier 4 arms the crisis gate: the engagement's roadmap pauses for a documented resolve-or-restructure conversation.

Tier 3 — Sponsor Intervention Required. A Governance Score below 60, any structural indicator active, or two Visibility & Contract indicators active. The pattern now requires the sponsor, by name, in a scheduled conversation with the register on the table.

Tier 2 — Friction. A Governance Score below 80, any ×2 indicator active, more than two approved scope changes inside thirty days, or an engagement still at Observer past day fourteen. Nothing is on fire; something is off pattern. Tier 2's obligation is documentation and watch — the register entry you make now is the receipt you'll need later.

Tier 1 — Normal Operations. The register is quiet. The obligation is to keep scoring honestly.

Tier assignment is computed, and the practitioner may override it — with a written rationale, and with both the computed and declared tiers journaled. The override exists because judgment outranks arithmetic; the journal exists because judgment must leave a trail.

5 · The renewal decision: windows, triggers, vocabulary

Renewals are lost in the last ninety days and blamed on the last meeting. The Standard's position: **the renewal is a process with a clock, not an event with a mood.**

The clock. Every engagement records its end date at signing, and a **renewal trigger** — a number of days before that end date at which the renewal process formally opens, whether or not anyone feels like opening it. When the window opens, the Standard expects three artifacts on the table: the Health Composite trend (Section 3), the current governance register and tier (Section 4), and the value ledger (Section 6). The renewal conversation is those three artifacts, narrated.

The vocabulary. Every renewal closes in exactly one of five recorded outcomes:

- **Extended** — same structure, new term.
- **Renegotiated** — new structure — scope, tier, or rate — new term.
- **Transitioned** — a planned, orderly handoff. Not a loss; frequently the correct outcome, executed well.
- **Wind-down** — a managed conclusion without succession.
- **Lost** — unplanned, unmanaged, or contested. The outcome the whole Standard exists to make rare.

The vocabulary is closed on purpose. "It's complicated" is not an outcome, and a practice that cannot classify its endings cannot compute its **renewal rate (M15)** — extended and renegotiated, over everything decided, trailing six months. That number is the practice's single most honest performance figure, and it is only as honest as the vocabulary is strict.

The doctrine. A renewal conversation that begins at the trigger, with the three artifacts current, is a business review. The same conversation begun at T-minus-three-weeks, from memory, is a plea.

6 · Operational Alpha: prove it in their numbers

Every fractional executive believes they generate more value than they cost. The Standard requires the belief to be a ledger.

The value ledger. Each entry records a delivered, quantified item of value — a cost eliminated, a margin recovered, a contract saved, a process whose cycle time you can price — together with its **implementation friction**: the real cost the client bore to realize it. Value claimed gross is marketing; value net of friction is evidence.

Net Alpha (M6). For any engagement:

M6 = Σ(entry value) − Σ(entry friction) − cost basis, where cost basis = retainer × months engaged.

Net Alpha is the number that answers the only question the renewal actually turns on: *net of everything, what has this engagement returned above what it cost?* When it is strongly positive, the renewal conversation is arithmetic. When it isn't, better to be the one who says so first, with the ledger open.

The ROI Bank (M7). The sum of Net Alpha across active engagements — the practice's cumulative, defensible, net delivered value. It is the number behind every rate increase you will ever propose.

The discipline. Ledger entries are recorded when the value lands, not reconstructed the week before renewal. Reconstructed value reads as reconstructed value; a ledger with dates reads as an operating record. And the friction column is not optional — including the client's cost to implement is precisely what makes the rest of the ledger believable.

7 · Running the Standard: cadence, tiers, and the practice layer

A standard you consult quarterly is a bookmark. FOS is run on a rhythm.

Weekly — the brief. One client-facing brief per engagement per week: what moved, what's blocked, what you need, with the current health snapshot. One rule elevates it from status mail to instrument: **the snapshot freezes when the brief is sent.** What you reported is preserved as reported — next week's numbers cannot quietly rewrite last week's story. A year of frozen briefs is an engagement's flight recorder, and at renewal it is unanswerable.

Continuous — the register. Governance indicators (Section 4) are marked as events occur, not batched at month-end. The instrument's value is its latency.

At the marks — the audit. The six-dimension audit at baseline, month three, month six (Section 3). Trend over level.

Quarterly — the business review. The Health Composite trend, the tier history, the ledger, and the roadmap, assembled for the sponsor. The QBR is the renewal conversation rehearsed three times before it counts.

Engagement tiers. The Standard contracts engagements at one of three intensity tiers, each with a cognitive-load multiplier the capacity math consumes:

- **T1 · Advisory** (×1.0) — direction, review, senior counsel; execution stays internal.

- **T2 · Active Leadership** (×2.0) — line ownership of designated functions, KPI accountability inside the contracted scope.
- **T3 · Transformation** (×3.5) — full operational command through a defined change mandate. *(Distinct from the governance tiers of Section 4 — one prices the engagement, the other polices it.)*

The practice layer. Four numbers make the portfolio legible:

- **Load (M11)** = Σ(engaged hours × tier multiplier) — hours weighted by what the hours cost you cognitively.
- **Utilization (M12)** = M11 ÷ (hour ceiling − admin reserve) × 100. Bands: under 75%, healthy headroom; 75–89%, approaching ceiling; **90% and above is a no-new-engagement condition**; above 100%, you are already paying for it.
- **Concentration (M10)** = largest retainer ÷ total retainers × 100. **Above 40% is the Single-Client Dependency failure mode**, named in Section 1.
- **Practice Score (M18)** = (avg M1 × 0.35) + ((100 − M10) × 0.25) + (min(M12, 100) × 0.25) + (avg M17 × 0.15) — engagement health, diversification, sustainable load, and governance, blended into the practice's one number.

M18's weights are an argument in miniature: engagement health carries the most because it is the product; concentration and capacity carry equal weight because either can kill a healthy book; governance is the leading edge that moves first.

8 · Implementation

Everything in this Standard can be run from a spreadsheet and a calendar. The registry of Section 2 is columns; the audit of Section 3 is a scoring tab and the discipline to write evidence; the register of Section 4 is ten rows, a date column, and the arithmetic of M17; the ledger of Section 6 is entries with a friction column. The Standard is doctrine, not software, and a practitioner who runs it manually is running it fully.

A full reference implementation conforming to FOS v1.0 is maintained as **FEOS — the Fractional Executive Operating System**: the eighteen canonical inputs, every metric, threshold, and tier rule in this document, the frozen-brief mechanism, the governance register, and the renewal workflow, implemented end to end in a single encrypted local file that keeps engagement records on the practitioner's machine and nowhere else. Conformance is verifiable by construction: any figure this document defines can be reproduced from the implementation's stored inputs.

Author's note

This Standard was engineered, not remembered. It began as a design problem: a failure pattern every practitioner recognizes and almost none can measure — authority leaving an engagement in increments too small to name. Solving it required the designer's discipline rather than the memoirist's: isolate the pattern from any single practice, weight its signals by structure instead of by anecdote, and refuse to publish until every number runs end to end in a working system. That is the test this document was held to. Nothing in it is aspirational; if the Standard says a number is computable, the reference implementation computes it.

The doctrine is free because a standard that hides is just a manual. Its authority rests on the only credential that survives contact with a skeptical operator: it is checkable. Score one engagement against the ten indicators. If the Governance Score doesn't match what your gut already knew, I want to hear about it.

— Raine Townsend