

METHODOLOGY WHITE PAPER

The Remidi™ Commercial Health Framework™

A Methodology White Paper for Investors and Operators

Updated: May 2026

Purpose

40% of PE-owned companies miss their value creation plans.¹ The critical questions are why, and what to do about it.

Leakage is commonly attributed to delayed resource deployment, slow decision-making, and missed execution against the plan. Those are accurate symptoms that all point to the same root cause. The commercial strategy layer that sits across the full revenue motion is largely unmeasured.

Marketing has HubSpot. Sales has Salesforce and Gong. Customer success has Gainsight. The commercial strategy layer (messaging, pricing, buyer experience, trust signals) has no equivalent.

Remidi developed the Commercial Health Framework to fill that gap. It is an unbiased, objective framework that can be applied consistently across companies. Actionable, benchmarkable, and trackable over time.

It is designed to be a resource investors and operators can return to across the full portfolio company hold lifecycle. From diligence, through the first 100 days, through value creation, and into exit preparation.

This document describes the framework, how scoring works, how the model has been validated, and how the data is used to inform decisions across the hold period.

Measurement Scope

Commercial strategy is the connective tissue across the full revenue motion. It guides demand generation, sales execution, customer success, and retention. Choices made at the commercial strategy layer (how the company is positioned, how it is priced, what evidence the buyer encounters, how easy it is to evaluate and adopt) shape the performance of every downstream function.

Remidi does not measure each of those functions directly. The framework focuses on the commercial strategy layer itself, observed as a prospective buyer would experience it. That layer informs the effectiveness of demand gen, sales, and CS even when it is not their direct measurement target.

The Remidi Commercial Health Framework evaluates what a prospective buyer experiences when they encounter the company: its positioning, pricing communication, proof points, conversion mechanics, and overall credibility. This is the space where a prospective buyer forms an opinion about whether a company is worth their time, before they ever take a sales call.

It is the layer operators can feel but cannot easily measure.

The framework complements the demand-gen analytics and CRM data operators and investors already see. It is additive, not duplicative.

What it does **not** measure: the company's internal people, processes, organizational design, sales team execution, or culture. Those things matter. They require deeper primary research, including interviews, observation, and internal data access, that public sources cannot reveal. Remidi can scope that work separately when it is warranted. The Commercial Health Framework is intentionally bounded to what the customer experiences, because that is what can be benchmarked at scale.

The Framework: Five Dimensions, 28 Sub-Dimensions

The framework scores each company across five dimensions.

DIMENSION	WHAT IT ASKS	SUB-DIMENSIONS
Messaging & Positioning	Is the story clear, differentiated, and proven?	Headline Clarity, Target Buyer Named, Quantified Outcomes, Value Lever Depth, Hard Numbers
Pricing & Packaging	Is the company capturing the value it creates?	Pricing Transparency, Tier Structure, Value Metric Alignment, Upgrade Path Clarity, Entry Point Accessibility
Buyer Experience	Can the buyer learn, evaluate, and self-qualify with low friction?	Try Before Buy, FAQ & Objection Handling, Setup & Integration Clarity, Low Commitment Terms
Trust & Credibility	Would a skeptical CFO or procurement team approve this?	Customer Logos, Testimonials, Case Studies, Review Platforms, Third-Party Validation, Security & Compliance, Company Stability, Support Responsiveness, Data & Privacy
Competitive Position <i>(data only)</i>	How clearly does the company differentiate against alternatives?	Differentiation Claims, Differentiation Evidence, Value Delta Articulation, Comparison Content, Win Narrative

Proprietary weightings and formulas are used to combine them into an overall score. The model has been calibrated through multiple iterations to maximize predictive validity against winner outcomes. The dimension and sub-dimension names are shared here so users can see the architecture of the model. The weights and formula stay

inside Remidi.

Competitive Position is scored on every company but does not contribute to the overall score. It serves as a control dimension for ongoing model validation.

How a Company Gets Scored

Deep research across the full web presence. Each scan includes a complete review of the company’s website along with its third-party footprint: G2, Capterra, TrustRadius, Gartner Peer Insights, LinkedIn, public funding and growth data, and category context. The same depth of review is then applied to the company’s direct competitors, in order to evaluate against the right comparable set.

Multi-LLM cross-validation. The framework is AI-native. Each scan runs through more than one large language model against the structured rubric. Specific rules dictate how to resolve disagreements between models. Disagreements that cannot be resolved are escalated for human review.

Structured, objective rubric. Each of the 28 sub-dimensions has a defined evidence threshold for each scoring level. The criteria are purposefully objective rather than subjective, to remove any judgment from the LLM that would introduce unwanted variability. A sub-score is determined by whether specific, observable evidence is present on the company’s owned and third-party surfaces, not by the analyst’s overall impression. The same rubric, the same evidence standard, every company. Two companies with identical scores on a given dimension are saying something genuinely comparable about that dimension.

Expert human review. Every score is reviewed by an experienced analyst before it lands in the database. The review catches context the LLMs miss: a rebrand in flight, a recent acquisition, a category quirk, an industry override on pricing transparency. Disagreements between the rubric output and the analyst trigger a second pass. The rubric drives the score. The expert ensures the score is defensible.

Rolling 6-month rescoring. Companies in the database are rescored on a rolling 6-month cadence. This serves two purposes. First, every portfolio company has a current score, not a stale one. Second, the database itself becomes a longitudinal asset that surfaces commercial strategy trends across the market over time, not just point-in-time benchmarks.

How Remidi Fits With Existing Tools and Approaches

Remidi is designed to complement and improve work and processes already in place. Funds and operators currently rely on three common approaches to assess commercial strategy. Each has trade-offs that Remidi was designed to address.

An **internal framework** applied by an in-house team gives the fund consistency across its own portfolio. Remidi adds competitor benchmarking from an independent third party, so the fund can see how its companies score against the broader market and against direct comparables, not just against each other.

A **custom diagnostic** from a consulting firm produces a thoughtful read on one company. Because each engagement is shaped to that company, the output is hard to compare to another engagement and hard to track once the project ends. Remidi allows customers to evaluate a plan's effectiveness on a common set of validated metrics, and to track ongoing impact. Remidi also produces specific, data-driven action steps including immediate quick hits.

An **LLM-driven ad-hoc analysis** through ChatGPT or a similar tool is fast, cheap, and generally insightful. Remidi adds consistency. Run the same assessment twice on the same company and the score is reproducible. That is not the case with an ad-hoc LLM prompt, and the variance is amplified when working across multiple LLMs.

The combination, applied consistently, is what produces a measurement system rather than a series of opinions.

Validation: Does the Score Predict Outcomes?

Remidi evaluates dozens of variables that, individually, are reasonable indicators of commercial strategy quality. The harder question is whether the combined score actually predicts which companies win in their categories. The validation work was designed to answer that question directly.

Methodology

The study population is the subset of the database that meets two conditions: the company is comparable to the broader middle-market cohort (filtering out outliers, non-SaaS models, and companies without enough public surface area to score), and the company has been included in the validation study (filtering out scans that fall outside the rubric's intended use). As of May 5, 2026, the validation cohort contains 648 companies.

Winner status is defined independently of the Remidi score. A company is classified as a commercial winner if it holds G2 Leader badge status combined with review volume and recency thresholds that indicate sustained category leadership in the most recent period. Winner determination uses only third-party signals. There is no path by which the Remidi score can influence winner status. That separation is what makes the test meaningful. In the current cohort, 127 of 648 companies qualify as winners (19.6%).

The test is whether the Remidi score discriminates winners from non-winners. The two metrics reported below address two different questions:

- **Predictability** asks how reliably the score sorts winners from non-winners across the score range.
- **Impact** asks how much the winner rate actually changes as scores move.

Both are reported because both matter. A model can be predictive in aggregate without producing meaningful effect sizes at the band level, and a model can show large band-level differences that turn out to be driven by a few extreme outliers. The validation work checks both.

Predictability

AUC = 0.773. AUC is a standard measure of how well a score separates two groups. In this case, winners from non-winners. A score of 0.5 is a coin flip. A score of 1.0 is perfect separation. A score of 0.773 is strong. The framework is doing real predictive work, not pattern-matching after the fact.

Validation is what makes the framework defensible to a board, an LP, or a portfolio CEO.

Impact

The score is not just statistically predictive. The size of the effect at the band level is large enough to be commercially meaningful.

The 18x finding. Companies scoring 66 or above on the Remidi Commercial Health Framework are approximately **18 times more likely to be commercial winners** than companies scoring 45 or below.

Winner rate by tier band, current cohort:

TIER	SCORE RANGE	WINNER RATE
Lagger	0–29	0%
Needs Attention	30–45	3%
Good Shape	46–65	14%
Best in Class	66+	39%

The winner rate rises as the score rises, with the steepest gains concentrated in the 60s and 70s. Modest score improvements at the inflection points translate into meaningful changes in winner odds. A company moving from the high-40s into the low-60s, or from the mid-60s into the mid-70s, sees a winner-odds shift that is much larger than the score-point shift would suggest.

That is the practical argument for treating Commercial Health as a measurement to manage, not just a snapshot.

The Benchmark Database

As of May 5, 2026, the database contains 889 scored companies, 655 of which meet the criteria for the middle-market benchmark cohort: meaningful revenue, clear B2B tech profile, beyond seed stage, below mega-cap. The database grows every day.

The middle-market subset is the relevant comparison set for an operator or investor. Scoring a portfolio company against the full universe would dilute the value of the data to our primary customers, investors and operators in the middle market.

New companies are added through scoring engagements and ongoing market research. The benchmark distribution is recalibrated as the database grows, so a score from this quarter is comparable to a score from last quarter.

When an operator or investor sees a portfolio company score a 47, the score sits in a defined distribution of comparable companies the reader would recognize.

Score Context: What the Numbers Mean

The framework places companies into four tier bands by overall score. Each band maps to a consistent set of recommended actions. The summary below is how to read a Remidi report.

External tier labels (used in client deliverables):

SCORE RANGE	TIER	WHAT IT INDICATES
0–29	Lagger	Bottom of the field. Systemic commercial gaps across multiple dimensions.
30–45	Needs Attention	Below market average. Dimension-level gaps require focused work.
46–65	Good Shape	Around or above market average. Specific dimensions worth optimizing.
66+	Best in Class	Top tier. Statistically associated with winner outcomes.

Universe averages vs. winners (current benchmark cohort):

DIMENSION	UNIVERSE AVG	WINNER AVG	GAP
Overall Score	56	71	+15
Messaging & Positioning	69	80	+11

DIMENSION	UNIVERSE AVG	WINNER AVG	GAP
Pricing & Packaging	42	61	+19
Buyer Experience	55	70	+15
Trust & Credibility	54	70	+16

The Pricing & Packaging gap is the largest. It is also the dimension where the universe is weakest in absolute terms. That combination has been one of our more actionable findings.

Common Questions

Is this just a website audit?

No. A website audit grades a website. The Remidi Commercial Health Framework grades the commercial posture a company presents to a buyer. The website is one input among several. The score reflects messaging, pricing, buyer experience, and credibility, measured against a benchmark cohort and validated against commercial outcomes.

Can you score companies with limited public footprints?

Yes, with caveats. Companies with thin public presences score lower on dimensions like Trust & Credibility, because that is what the buyer experiences. If the company has invested in commercial visibility outside the public domain (partner channels, direct sales, embedded distribution), the human review step flags this and the findings note the boundary.

What if the score conflicts with what we already know internally?

A company that performs well internally but scores low externally has a translation problem. The commercial reality is not reaching the buyer. A company that scores high externally but underperforms internally has a different problem worth investigating. The score is a measurement, not a verdict.

How fresh is the benchmark?

The database is rescored on a rolling 6-month cadence. The benchmark distribution is recalibrated as new companies are added and as existing companies move. A score from this quarter is comparable to a score from last quarter and to scores from prior periods.

Closing Note

The Remidi Commercial Health Framework gives investors and operators a common language and a common measurement for the commercial layer of a portfolio company. It is not a strategy. It is not a substitute for operator judgment. It is a measurement system, built to be unbiased, applied consistently, benchmarked against comparable companies, and validated against commercial outcomes.

Used well, it surfaces what to fix first, makes progress visible, and makes the portfolio comparable to itself and to the market.

That is the entire objective.

For a methodology consultation or to discuss applying the framework to a specific portfolio, contact Remidi at remidiworks.com.

FOOTNOTES

1. Reported by Alvarez & Marsal (2025), with consistent findings from Simon-Kucher and EY.