ValueLogics.ai

THE ROI AGENT COMPANY

Sarnosites

PROPOSAL · SYSTEMS INTEGRATOR · LEGACY MODERNIZATION

Renewal Rescoping

The business case is the only document in an enterprise that is never audited. Two artifacts change that — and change what a renewal is argued on.

PREPARED FOR BANGALORE SOFTSELL

3	4	2
INSTRUMENTS	PROOF TESTS PER CLAIM LINE	PARADOXES — ONE OF THEM OURS

> **What this document is.** A proposal from **ValueLogics.ai** to **Bangalore Softsell** for a single-account pilot of **Renewal Rescoping** — rebuilding one upcoming modernization renewal as a documented value case instead of a scope-and-rate negotiation. It is written to be read by the person who signs the renewal on the customer side as much as by the person who defends it on yours.

The renewal problem nobody names

A legacy modernization program is sold on a future state. It is renewed on an invoice.

Somewhere between those two conversations, the argument changes hands. Phase 1 was sold to a sponsor who wanted off an unsupported platform. Phase 2 is renewed by a procurement team comparing your blended rate to three other SIs, with a spreadsheet that has a column for cost and no column for anything else. Your account team answers the only way it can — with effort. Headcount. Velocity. Percentage of the estate converted. Milestones met.

Effort is not value. It is the **price** *of value. When the only number on the table is the price, the negotiation has exactly one direction to go.*

The Paradox of Excellence

The paradox was named twenty years ago and it is still with us. It is worth being precise about why, because the answer is not that the diagnosis was wrong.

> **The diagnosis, and the step where it stalls.** Mosby and Weissman named the Paradox of Excellence in 2005: perform consistently well and customers stop noticing, until excellence is taken for granted and the relationship is repriced as a commodity. Their remedy is the Continuous Visibility Wheel — discover the expectations, define your distinguishing value, select the metrics to render visible, **uncover the best source of data to use**, then deliver the information. The diagnosis has held up for two decades. Implementations stall, almost without exception, at the fourth step.

They also ask three questions in a row about value, and the order is the whole story: do your customers *know* your distinguishing value, can they *articulate* it, can they *quantify* it. Most vendors reach the third question and stop. Quantification is not a harder version of articulation — it is a different kind of thing. It needs a frozen baseline and a matched actual, and those are artifacts, not talking points.

Here is why the fourth step stalls, and it is the part almost no SI has articulated to a customer:

When you retire a legacy subsystem, the cost of running it **stops appearing in the customer's ledger**. The MIPS charge falls off. The vendor support contract lapses and is not renewed. The two contractors who knew the RPG are released. The DR capacity is reclaimed. None of this is recorded as a saving — it is recorded as an absence.

THE PARADOX OF EXCELLENCE, IN THE LEDGER

The act that creates the value is the act that erases the evidence for it

These are not two events, one unfortunately following the other. They are the same event. You demonstrate your worth by removing a cost — and a removed cost leaves no line in the ledger to point at. The cleaner the removal, the more complete the erasure.

Within one budget cycle the customer's baseline has quietly absorbed everything you delivered, and the next planning round starts from the improved number as though it had always been true. At renewal you are then asked to justify the next term against a baseline that already contains your own results, and to do it without being permitted to count them. **Every successful phase weakens the next renewal.**

This is not a selling problem, and it is not a communication problem either. Telling the customer more loudly about a saving that has no line in their ledger does not put a line in their ledger. It is a **measurement** problem that shows up at selling time, and it has exactly one exit: the evidence has to be written down before it is erased, because afterwards there is nothing left to write it down from.

The wheel was never wrong. It was missing an instrument. Steps one to three are strategy and can be done in a room; step four is forensics and cannot.

> **The window, and why it closes.** Erased is not quite the same as gone. Contracts, invoices, change logs and capacity records hold the record for a few years — usually a little longer than the people who remember what it was like before. After that the erasure is complete and the term is simply unprovable. A replay run today on a three-year-old phase is archaeology with a deadline; the same replay in two years is not available at any price.

What the ROI Framework actually holds

Most ROI conversations in modernization are half a conversation. They run either as a TCO case — here is what the old estate costs, here is what the new one costs — or as a benefits case — here is the agility you will gain. Procurement trusts the first and discounts the second, so the benefits half gets stripped out and what is left is a cost comparison you can only win by being cheaper.

The ValueLogics ROI Framework carries both sides in one ledger, against one baseline, in the customer's own measures. TCO is not a separate document from benefits; it is the cost face of the same value map. That is what makes the case survive procurement: every benefit line can be traced back to a cost line or a measured quantity that the customer's own finance function already reports.

The four buckets, applied to modernization

- **Cost savings — run cost eliminated.** Platform licences, MIPS or capacity charges, vendor support and maintenance contracts, scarce-skill contractors, the maintenance window itself.
- **Asset savings — capital avoided.** Hardware refresh deferred or cancelled, DR footprint reclaimed, environments collapsed, data-centre floor released.
- **Revenue — throughput restored.** Change requests delivered per quarter, time from approved requirement to production, product or rate changes that could not previously be made at all.
- **Risk reduction — exposure retired.** Key-person dependency on retiring skills, unsupported-platform audit exceptions, regulatory findings, failed-change incident rate.

Shape of a typical modernization ledger.
Your account's actual split is measured, not assumed.

The document that is never audited

Every enterprise software engagement begins with a business case. Somebody built a model, put a number on it, and used that number to get the deal approved.

Then the deal is signed, and that document is never opened again.

It is not filed. It is not tracked. Nobody goes back at the end of the term and asks whether the number was right. It is the only financial document in a large organisation that is produced to justify a decision and then permanently excused from ever being checked. Procurement knows this. That is precisely why the benefits half of your case gets discounted to zero before the conversation starts — not because the benefits were implausible, but because everybody in the room has learned that nobody ever comes back and settles up.

The business case is treated as sales collateral because it has always behaved like sales collateral. Change that behaviour and the whole document changes weight.

The second paradox

The first paradox is not ours. The second one is, and it is the reason the first is solvable at all.

THE DOUBLE PARADOX

One — the act that creates the value erases the evidence for it.

Two — the only number worth having is one that could have gone against you.

A measurement you control is worth exactly as little as your control over it. A vendor's own ROI figure is discounted on sight, and the discount is rational: everyone in the room knows the method was chosen after the answer was known. Confident language does not fix this. Nothing said by the party with the interest fixes this.

So the measurement runs regardless of sign. A Replay that comes back under the claim is published as a Replay that came back under the claim, in the same format, on the same four tests. That is not a disclaimer attached to the product. It is the product. The willingness to print a shortfall is the entire reason a surplus means anything.

Together the two paradoxes say: you must prove it, and you must be willing to be disproved. Which raises the obvious commercial objection — why would any vendor build an instrument capable of convicting them?

Why that is survivable: ROI DriftWatch

Because value does not fail at renewal. It fails quietly in month seven.

A Replay run on its own is a bet placed at the end of the term, on evidence nobody was watching while it accumulated. **ROI DriftWatch is the same measurement running continuously against the design intent** — the value the contract said the system would produce. Divergence surfaces while it is still divergence and not yet outcome.

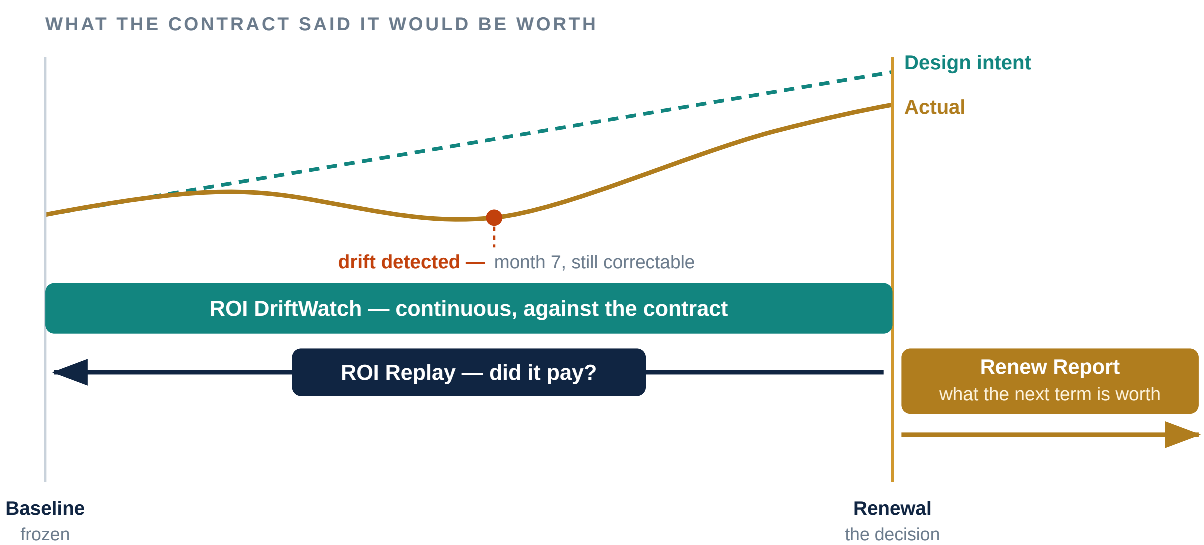

One term, three instruments. Replay is the audit. DriftWatch is the reason the audit rarely has to deliver bad news.

That is what it means to say the designed outcome is assured, and it is worth being exact about the claim. Not that nothing goes wrong. That nothing goes wrong *invisibly*. A gap found in month seven costs a conversation and a correction. The same gap found at renewal costs the renewal.

> **The two commitments are one commitment.** Reporting regardless of sign, without watching continuously, is reckless — you have agreed to publish a number you have no way to influence. Watching continuously, without committing to report regardless of sign, is theatre — you are collecting evidence you intend to use selectively. Each is unwise alone. Together they are the only arrangement in which a vendor can afford to be audited, and the only one in which being audited is worth anything.

Three instruments

One looks backward, one runs throughout, one looks forward. Between them they are the fourth step of the wheel, built rather than recommended.

INSTRUMENT ONE

ROI Replay™

The original business case, re-run against what actually happened. Claimed versus delivered, line by line, against a baseline frozen from the customer's own pre-engagement records. Every line carries its source, its method, its attribution note and a confidence grade — and lines that cannot carry all four are struck, not softened.

A replay is an audit of a document nobody audits. It is the answer to *did the last term actually pay*, produced from evidence rather than from the account team's recollection.

The specimen below is **shape only** — it shows the columns, not your account's numbers. Those are an output of the work, and some of them will come back lower than the claim:

CLAIM LINE	CLAIMED	DELIVERED	VARIANCE	GRADE
Platform support contract retired	—	—	—	A
Specialist contractor spend released	—	—	—	A
Change requests delivered per quarter	—	—	—	B
Hardware refresh avoided	—	—	—	A
Audit exceptions closed	—	—	—	C
Earned to date	—	—	—	—

> **Why the downside is the point.** A replay that can only come back favourable is a brochure. The reason a replayed number carries weight in a procurement conversation is that the same method could have returned a shortfall and would have printed it. That is the whole of the credibility, and it cannot be simulated by confident language.

INSTRUMENT TWO

ROI DriftWatch™

The claim, held against the work as the work happens. Measures captured at the moment the activity that moves them completes, scored on the same four tests, compared to the design intent the contract carries. It answers *is it still paying* every month rather than once at the end — and it is what turns the next Replay from archaeology into an export.

INSTRUMENT THREE

Renew Report™

What is still on the table, and what the next term is worth. Earned-to-date carried forward from the Replay; baseline drift made explicit; the remaining gap stated bucket by bucket as a quantity of unrealised value rather than a list of subsystems; and the proposed term priced against that gap.

The Renew Report is what converts the renewal from *how many people for how many months* into *how much is left and what does closing it require*. Those are two different negotiations, conducted by two different people, and only one of them is decided on rate.

Baseline drift is the line that does not exist anywhere else. It is the quantity your work already produced that the customer's budget silently absorbed — named, dated and restored to the ledger. Without it the customer is comparing your next term against a baseline that contains your own results. With it, the comparison is honest for the first time.

Where they live

None of the three is a deck. All are views of the same value map your delivery work already runs against — the same data, read by different people at different moments.

CONTEXT	WHO IS LOOKING	WHAT THE VIEW HOLDS
Sell	Prospect, pre-contract	The claim, against an assumed baseline
Deliver	Delivery team, in term	ROI DriftWatch, against design intent
Renew	The signer, at decision	ROI Replay, then the Renew Report

Nobody buys a decision support system. They buy a defensible answer at a moment when being wrong is expensive. The system is how the answer gets produced a second time without starting over.

The proof standard

Four tests. Every line in a Replay is scored against them before it is allowed into the document, and the grade is **printed on the line** — so the customer's finance function can audit the case without asking you anything.

- **Baseline.** A recorded pre-state, dated, from the customer's own system.
- **Attribution.** A stated reason the change is attributable, and a stated alternative explanation considered.
- **Matched actual.** A post-state measured the same way, on the
- **Our units.** Expressed in a unit the customer already reports.

same definition.

A percentage with no source scores zero and does not appear. Most modernization business cases in circulation today would fail this rule on the majority of their lines.

Why nobody looks

Everything above assumes the obstacle is cost, effort, or missing data. It is none of those. The evidence can usually be recovered, the reconstruction is bounded, and the method is not difficult. The obstacle is that nobody wants to know.

Ask an executive to replay a three-year-old business case and watch the objections arrive in order: the data is messy, attribution is hard, it was a different market, the sponsor has left. Every one of those is true and every one is surmountable, and none of them is the reason. The reason is that everybody in the room already suspects the case did not pay, and looking is what makes the suspicion official. Not looking is individually rational for every single person in the building. That is exactly why it is universal.

THE FEAR THAT KEEPS THE PARADOX ALIVE

Measurement is treated as an autopsy.
It is a steering input.

The fear assumes the number is already fixed and that measuring merely uncovers it. That is true of one thing only — an audit run at the end of a term — and false of everything before it.

A programme measured from month one has as many available outcomes as it has months left in which to intervene. A programme nobody measures has exactly one: whatever happens. The unmeasured case is not the safe case. It is the case whose distribution has already collapsed to a single point that no one has read yet.

Looking does not reveal the number. Looking changes it.

That is the argument for DriftWatch over a terminal audit, and it is the honest answer to *what if it comes back bad*: then you found out while it was still a decision rather than a result. The organisations that most fear the replay are the ones for whom watching would have been worth the most, and they are the same organisations, for the same reason.

The alternative is not neutral. Turning a blind eye is itself a choice, and it has a destination. An organisation that never proves value can only judge price. A vendor that is never credited with value can only compete on price. Each side reinforces the other, and the end state is a market in which cost is the only visible variable and the only remaining move is to be cheaper than you were last year.

WHERE NOT LOOKING LEADS

The race to zero is not something that happens to a market.
It is what a market does in the absence of measurement.

Which leaves one inversion worth stating plainly. Measuring to a renewal does not only defend a renewal — it re-points the work at something other than effort, for the straightforward reason that people deliver against whatever is going to be checked. Alignment is downstream of measurement, not upstream of it.

Value-aligned organisations are not built by declaring values. They are built by attaching a number to a date.

What is actually being bought

Not the number. The number is consumed once and discarded, like every number before it. A figure that wins one renewal has no residual value and cannot be sold twice.

What lasts is the score on the evidence. A Replay does not only report what the last term earned — it reports **how much of what was claimed turned out to be provable at all, and to what standard**. Those are two different findings, and the second one is worth more than the first.

THE REAL INSTRUMENT

We do not measure the value.

We measure the measuring.

Every claim line carries a grade. Aggregate the grades and you no longer have a statement about one vendor's performance. You have a statement about an organisation's capacity to know whether anything it buys works.

When sixty per cent of a claim set scores zero, that is not a verdict on the vendor. It is a finding about the buyer, and it holds whichever vendor is examined next — which is exactly why it is worth having. An organisation that cannot grade its own evidence is not choosing between vendors. It is guessing, expensively, on a schedule.

The four tests are the product. A Replay is one application of them.

That has a consequence worth stating plainly rather than discovering later. An organisation that adopts the tests eventually stops needing anyone to run them and starts needing someone to *certify that they were run*. That is a smaller engagement than a Replay, a considerably longer one, and the only one that scales past the capacity of the people doing the reconstruction.

What this is worth beyond the renewal

An ROI Replay is a document your customers would pay for on their own account, about their own vendors — and a document your competitors cannot answer, because answering it means opening a business case they have no evidence for.

We have built exactly this before, commissioned by a service provider, addressed to their customer, and used as that provider's primary sales instrument. It moves an SI from defending a rate to publishing a result.

The ask. One account with a renewal falling in the next two to three quarters. We run the Replay first — including the finding that the evidence is not there, if that is the finding. Nothing beyond it is committed until you have read it.

ValueLogics.ai · THE ROI AGENT COMPANY · Sarnosites

Author once. Generate anywhere. Prove it paid off. valuelogics.ai · Prepared September 2026