ValueLogics.ai

THE ROI AGENT COMPANY

Sarnosites

CHANNEL OFFERING · SALESFORCE · COMPLETES SF VALUEANALYTICS V2

ROI Replay for Salesforce

We do not compete with Salesforce — we complete it. The business case is the one document in an enterprise that is never audited. This offering audits it, on a date the calendar already carries.

STRATEGIC BRIEF

4	40%+	9
BUYING MOMENTS	AGENTIC PROJECTS CANCELLED BY 2027 — GARTNER	BUSINESS AREAS BECOME THE MEASURE LIBRARY

What this completes. SF ValueAnalytics v2 (nine business areas) was never wrong. It was unfinished. The nine business areas stop being the offering and become the **measure library the offering draws on** — and the step that was always missing gets added: settlement. Measurement without settlement is a dashboard. Measurement that settles is an audit.

THE CHANNEL POSITION, IN ONE LINE

We do not compete with Salesforce. We complete it.

Salesforce sells the platform and the transformation. Nobody comes back afterwards and settles up on what either one earned. That gap is not a weakness in the platform — it exists around every enterprise platform ever sold. It is simply unoccupied.

Why completion rather than replacement

ValueAnalytics sold a capability. Analytics across nine business areas — genuinely useful, genuinely hard to build, and impossible to attach to a moment. Nobody wakes up needing analytics. There is no date on the calendar when a capability becomes urgent, so the offering competed on interest rather than on necessity, and interest does not have a budget line.

ROI Replay sells to a moment. The renewal. The annual true-up. The AI business case coming up for its first review. The conversation that starts *why are we still paying for this*. Every one of those has a date, a person who owns it, and money already allocated to the decision.

A capability has to create its own urgency. A moment arrives on its own and brings a budget with it. The nine business areas were never the problem — they were the answer to a question nobody had scheduled.

The argument

Every Salesforce estate was implemented on a business case. Somebody built a model, put a number on it, and used that number to get the programme approved.

Then the deal was signed, and that document was never opened again. It is the only financial document in a large organisation produced to justify a decision and then permanently excused from ever being checked. Procurement knows this. It is exactly why the benefits half of any subsequent 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.

There is a second reason, and it is the one partners feel rather than name. Mosby and Weissman called it the Paradox of Excellence in 2005 — perform consistently well and customers stop noticing, until the relationship is repriced as a commodity. Their remedy, the Continuous Visibility Wheel, runs: discover the expectations, define your distinguishing value, select the metrics to render visible, **uncover the best source of data to use**, deliver the information. Twenty years on, implementations still stall at the fourth step — because consolidation work erases its own evidence. Retire a legacy CRM licence, collapse an integration, release the contractor who maintained it — and the cost stops appearing in the ledger. It is not 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

The same event, not two. You demonstrate worth by removing a cost, and a removed cost leaves no line to point at. Within a budget cycle the baseline has absorbed it, and the next funding round starts from the improved number as though it had always been true.

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 OFFERING IN ONE LINE

We settle up.

ROI Replay re-runs the original business case against what actually happened — claimed versus delivered, line by line, against a baseline frozen from the customer's own pre-implementation records. Renew Report carries the earned figure forward into what the next term is worth.

The second paradox, and why the offering survives it

The Paradox of Excellence 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 the vendor controls is worth exactly as little as that control. A partner'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. **So the measurement runs regardless of sign.** A Replay that comes back under the claim is published as one, in the same format, on the same four tests. That is not a disclaimer attached to the product. It is the product.

Which raises the commercial objection immediately: why would a partner commission an instrument capable of convicting them?

Because value does not fail at renewal — it fails quietly in month seven. ROI DriftWatch is the same measurement running continuously against the design intent the contract carries, so divergence surfaces while it is still divergence and not yet outcome. A gap found in month seven costs a conversation. The same gap found at renewal costs the renewal.

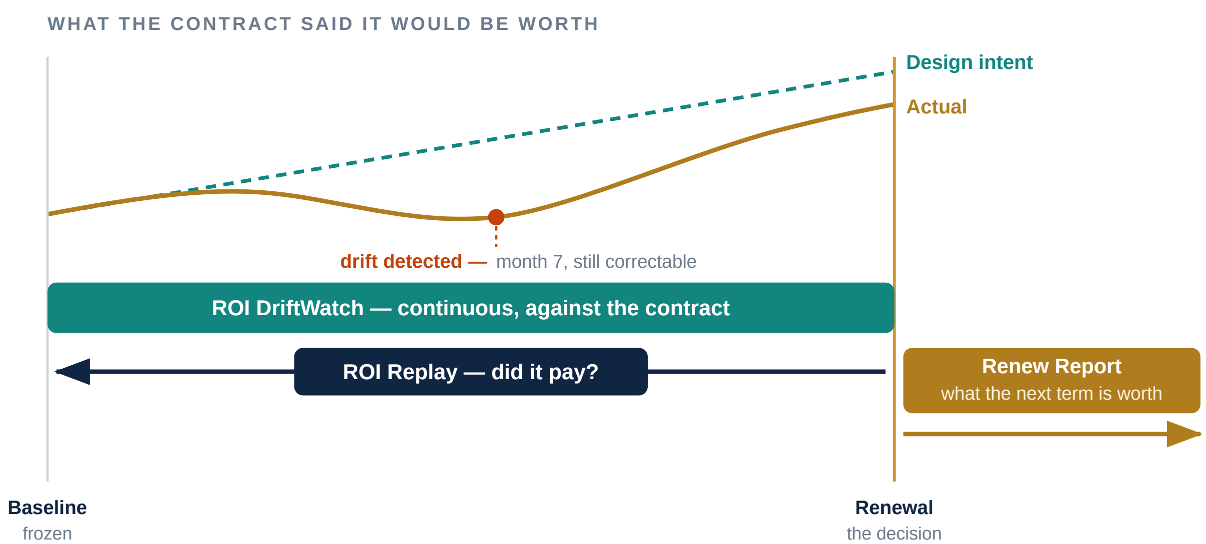

One term, three instruments. Replay is the audit. DriftWatch is the reason the audit rarely has to deliver bad news.

Reporting regardless of sign without watching continuously is reckless. Watching continuously without committing to report regardless of sign is *theatre. Together they are the only arrangement in which a partner can afford to be audited — and the only one in which being audited is worth anything.*

Four moments in a Salesforce estate

- **Renewal and true-up.** Annual, dated, owned by someone with signature authority. The question on the table is always the same one a Replay answers.
- **Partner SOW extension.** The SI or ISV defending its next statement of work against procurement, with effort as its only available argument.
- **Agentforce funding review.** The first wave of agentic work coming up for its second round of money.
- **Consolidation or platform review.** Somebody asking whether the estate earns what it costs, usually with a consultant already in the room.

The Agentforce case, specifically

Gartner expects **more than 40% of agentic AI projects to be cancelled by 2027**. The interesting question is not whether that number holds. It is *what a cancellation actually is*.

A cancellation is a funding decision made in the absence of evidence. Nobody cancels a programme that can show what the last tranche earned. They cancel the one where the honest answer to *did it pay* is a shrug and a slide about adoption. The first agentic wave is now old enough to be asked that question, and almost none of it can answer.

A Replay on the prior wave is the precondition for funding the next one. That makes this offering countercyclical: it sells hardest exactly when AI budgets tighten, because tightening is what forces the question.

Who buys it in the channel

- **The SI or ISV partner** buys a Replay to defend a renewal or SOW extension — moving from arguing a rate to publishing a result.
- **The customer** buys a Replay to decide — and increasingly to decide about a vendor who would rather they did not ask.
- **The partner who commissions one** about their own work is making a bet their competitors cannot match, because matching it means opening a business case they have no evidence for.
- **Co-brand only.** Partner authors the commitment; ValueLogics makes it reconcilable at runtime. The Replay *is* the reconciliation.

The proof standard

Four tests. Every line in a Replay is scored 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 anyone.

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

> **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 no amount of confident language substitutes for it.

Capacity

The first Replay on any account is reconstruction work — evidence has to be recovered from change logs, invoices, contracts and people's memories before anything can be scored. That is real, bounded, non-compressible effort, and it does not parallelise well.

So this books out. A limited number of Replays can be delivered before year end, and the honest way to launch it is to say so on the page rather than discover it at the third signature. Scarcity that is true is worth more than scarcity that is staged — and a waitlist with a date on it is a stronger commercial position than an open calendar.

The second Replay on the same account costs a fraction of the first, because the baseline is already frozen and the measures are already instrumented. That is the subscription, and it should be bought because the first one proved it worth buying.

The channel ask. One partner account with a renewal or SOW extension inside two to three quarters, and one Agentforce programme coming up for funding review. Those are the two shapes worth proving before the page goes live — one defends a vendor, one decides about one.