

THE POST-QUOTE INTELLIGENCE RESEARCH SERIES

The Post-Quote Intelligence Gap

Volume I

Category Definition, Evidence Review and Research Agenda

The foundational volume of the Post-Quote Intelligence Research Series, examining the period between quote delivery and customer decision.

Post-Quote Intelligence · a specialised application of · **Customer Decision Intelligence**

Izhak Salomonovitch

Founder, Zevik

Volume I · Published 2026 · Last reviewed 21 July 2026

zevik.io

Preface: A Research Series, Not a Single Paper

This document is Volume I of the Post-Quote Intelligence Research Series — the first stage of an ongoing enquiry into what happens in the period between a quote being delivered and a customer deciding. Its purpose is deliberate and bounded. Volume I defines the operational problem, reviews the evidence currently available, examines the limits of existing software categories, and proposes a working framework — the Five Levels of Customer Intent — for interpreting customer behaviour after a quote is sent. It also sets out an explicit research agenda: the propositions that later volumes will test.

It is stated plainly, and treated as a strength rather than an omission, that little rigorous public data exists about this phase. That is precisely why a foundational volume must come first. A category cannot be measured before it has been defined, and a survey cannot be designed before the questions worth asking have been identified. Volume I does that definitional work. Later volumes will carry the empirical load:

- Volume II — a structured survey of service businesses, testing how quotes are followed up in practice and where existing systems fall short.
- Volume III — behavioural validation, examining which observed signals actually correlate with customer decisions once sufficient anonymised data exists.
- Subsequent volumes — periodic industry benchmarks as the evidence base matures.

The reader should therefore approach Volume I as what it is: a category definition and a research agenda, offered as a framework to be tested, challenged, and refined by the evidence that follows — not as a set of settled empirical findings. Where this paper reasons rather than measures, it says so.

Contents

If this page appears blank, right-click the area below and choose "Update Field" to generate the contents.

Executive Summary

Service businesses invest heavily in two activities: acquiring the enquiry, and producing the quote. Both are measured, tooled, and optimised. Marketing spend is attributed. Quoting software has become fast and professional. Yet the period that sits between these two investments and the revenue they are meant to produce — the interval *after* a quote is delivered and *before* the customer decides — remains largely unmanaged. Parts of it are now recorded as isolated events, but it is rarely interpreted, seldom measured as a decision process, and in most businesses not turned into a clear next action.

This paper argues that this interval is a structural gap in the service-industry technology stack, not a matter of individual sales discipline. It is careful about the nature of that gap. Existing platforms increasingly capture individual post-quote events — a quote viewed, a PDF downloaded, a comment left, an acceptance or a decline. What remains uncommon is a dedicated capability that combines customer engagement, explanation consumption, questions and response actions into an **interpreted assessment of intent and a recommended next step**. Customer Relationship Management systems manage the relationship. Job-management and field-service platforms manage the work. Quoting and accounting tools generate the document and record its administrative status. What none of them was built to produce is an answer to the question every operator actually asks after hitting *send*: **what is this customer doing, and what should I do about it, today?**

We give this gap a name — **the Post-Quote Intelligence Gap** — and position the capability that closes it as a discipline in its own right. **Post-Quote Intelligence** is the application of a broader field, **Customer Decision Intelligence**, to the specific window between quote delivery and customer decision. The paper sets out the operational problem, explains why existing categories of software — though they increasingly capture isolated post-quote events — were not built to interpret them into a decision, examines the cost where credible data exists, and proposes a working framework, the Five Levels of Customer Intent, that businesses can use to reason about engagement after a quote is sent.

A deliberate discipline runs through this document: where reliable third-party data exists, it is cited and its provenance is stated plainly. Where it does not — and in this field it frequently does not — we say so, and make the argument through workflow analysis and logic rather than through invented figures. As the paper shows, the post-quote phase is so poorly instrumented that some of the most widely circulated “statistics” about sales follow-up turn out to have no verifiable source at all. That absence of evidence is itself part of the argument.

Many service businesses that issue priced proposals retain an under-managed phase in the sales process. The question this series asks is not merely whether it exists, but how much it costs, and what it would take to close it.

1. Defining Post-Quote Intelligence

Because this series aims to establish Post-Quote Intelligence as a recognised category, a precise definition comes before the argument, so that readers — and the automated systems that increasingly summarise documents like this one — can extract and reproduce it accurately.

DEFINITION — POST-QUOTE INTELLIGENCE

Post-Quote Intelligence is the systematic capture, interpretation and operational use of customer behaviour in the period between the delivery of a priced proposal and the customer's decision.

It is a specialised application of the broader discipline of **Customer Decision Intelligence** — the practice of making a customer's decision process visible and interpretable so that action follows evidence rather than assumption.

1.1 What Post-Quote Intelligence is not

The category is defined as much by its boundaries as by its content. Post-Quote Intelligence is distinct from, and does not reduce to, any of the following:

- **Basic email-open tracking.** A single open is a weak, noisy signal (see §4.2); interpretation across multiple behaviours is the point.
- **Quote creation.** Producing the priced document is the job of quoting and accounting tools; the category begins after the quote is sent.
- **CRM pipeline management.** Tracking a relationship through pipeline stages is not the same as reading a specific customer's decision on a specific open quote.
- **Fixed follow-up automation.** Reminders on a timer are not conditioned on the customer's actual behaviour; behaviour-conditioning is central here.
- **Online quote acceptance.** Recording that a customer clicked “accept” or “decline” is an administrative status event, not an interpreted assessment of intent while the decision is still open.
- **Generic sales analytics.** Aggregate dashboards describe a pipeline in the past tense; Post-Quote Intelligence is about the live, in-flight decision on an individual quote.

1.2 Scope of this paper

SCOPE

This paper focuses on service businesses that regularly issue priced proposals as part of their sales process, and where quote volume, customer complexity or team size make purely manual follow-up increasingly difficult to do consistently. The discussion is most directly applicable to trade, construction and field-service businesses — electrical, plumbing, HVAC, building, installation, maintenance and commercial-services contractors — although many of the principles may apply more broadly.

The category is defined by operating characteristics, not by a fixed headcount or revenue band: what matters is that a business sends enough priced proposals that following each one up well, by hand, becomes hard to sustain. While the examples in this series draw heavily on field-service and trade businesses, the concepts may extend to other proposal-driven industries and geographies wherever a priced offer is followed by a decision the seller cannot see.

2. The Hidden Revenue Leak in Service Businesses

Consider how a typical service business — a trade contractor, an installer, a small commercial services firm — spends money and effort to produce a single unit of revenue. Three distinct investments are made before a dollar is earned, and they are made in sequence.

2.1 Three investments, one blind spot

1. **Acquiring the lead.** The business pays for visibility — search advertising, directory listings, referral relationships, a van with signage, a website. Every enquiry has a real, often measurable, acquisition cost.
2. **Producing the quote.** Someone attends the site, scopes the work, prices it, and prepares a document. This is skilled labour and it is not free. For many operators it happens in the evening, after the paid work of the day is done.
3. **Waiting.** The quote is sent. And then — nothing structured happens. The business waits, guesses, and eventually either wins the job, loses it, or never finds out which.

The first two investments are treated as business activities to be managed. The third is treated as weather — something that happens to the business rather than something the business does. Yet it is precisely the stage at which the value of the first two investments is either realised or lost.

2.2 Why the leak is hidden

The leak is hidden for a structural reason: little *interprets* it, even where fragments are recorded. A lost lead that never became a quote shows up as a low conversion rate on marketing. A produced quote shows up in the quoting tool, and its status may later read “declined” or simply remain “sent.” But the *reason* a quote quietly died — the customer was confused, or comparing, or busy, and no one followed up at the right moment — is rarely captured anywhere, because the decision process in which that reason formed was not observed as a process. The quote is recorded, at best, as one with no linked job and perhaps a bare status. The *reason* it did not convert is never captured, because the period in which that reason formed was never observed.

This is what makes the post-quote phase a **revenue leak** rather than simply a low win rate. A low win rate is a known quantity you can work on. A leak is a loss you cannot see, cannot attribute, and therefore cannot fix. The money does not fall out of a place you are looking.

A note on evidence. Throughout this paper, “quote” is used to mean any priced offer a service business sends to a prospective customer — estimate, quotation, proposal, or bid. The dynamics described apply wherever a business delivers a price and then waits for a decision it cannot see.

3. The Quote Visibility Gap

Once a quote leaves the business, it commonly enters what many operators describe, in our conversations and in the trade’s own language, as a **black box**. As Section 4 details, some systems now record fragments of what happens next — a view here, an acceptance there — but for most operators those fragments do not add up to a picture. The business is left making decisions — to call or not to call, to hold price or discount, to pursue or move on — with little interpreted information about the one party whose behaviour determines the outcome.

3.1 The questions with no answers

In the days after a quote is sent, an operator is silently carrying a set of questions. A system might log that one or two of them occurred; what none reliably answers is what they mean or what to do about them:

- **Has the customer even opened it?** Some systems can now record an open — but the record is ambiguous and, as Section 4.2 shows, noisy. An open could mean the quote was glanced at and dismissed, or read three times and forwarded to a partner. The event may be visible; what it signifies is not.
- **Have they understood it?** A quote is a compressed technical and commercial document. A customer who does not understand a line item, a scope boundary, or a payment term may simply go quiet rather than ask. Confusion and disinterest look identical from the sender's side.
- **Are they comparing?** In many service categories a customer will hold more than one quote. The business typically has no view of where it sits in that comparison, or whether the comparison is even still live.
- **Have they simply become busy?** Some quotes do not lose so much as stall — the customer intended to proceed and life intervened. A well-timed nudge might recover the job; a mistimed one might irritate. Without visibility, the business cannot easily tell which case it is looking at.
- **Should I follow up today, next week, or not at all?** This is the operational question every other question resolves into. Absent an interpreted signal, it is often answered by guesswork, timing convention, and whichever quote happens to surface in memory rather than by evidence about the specific customer.

3.2 The cost of the black box is a decision, not just a feeling

It is tempting to treat post-quote uncertainty as an emotional problem — the low-grade stress of unanswered quotes accumulating in the back of an operator's mind. That stress is real. But the substantive cost is that **follow-up decisions are largely made without interpreted evidence**. The business is not necessarily choosing badly; it is often choosing blind. Some blind choices happen to be right. Many are simply uninformed, and the aggregate of uninformed choices across a year of quoting plausibly adds up to a meaningful amount of lost and mis-spent effort.

Silence is among the most expensive signals a service business receives — constant, and, without interpretation, almost impossible to read.

4. Why Current CRM and Quoting Platforms Don't Solve This

The natural objection is that this problem must already be solved — that some category of existing software surely covers it. The honest answer is more precise than a flat “no,” and the precision matters. Several established platforms now capture *individual events* from the post-quote window. HubSpot's quoting tools log buyer activity on a quote — sent, viewed, downloaded, printed, accepted and signed — and surface it on a timeline (with the caveat that specific availability depends on plan and quote type).^[4] Xero's online quotes let a customer accept, decline, and leave comments or questions through the online quote, and move it through Draft, Sent, Accepted, Declined and Invoiced statuses, notifying the sender.^[2] Field-service tools such as ServiceM8 similarly support online quote viewing and acceptance. So it is simply not true that existing systems see nothing after a quote is sent.

What these systems provide, however, is a set of **fragmented events and administrative statuses** — discrete facts (“viewed,” “declined,” “accepted”) recorded against a record. What none of them was designed to produce is the layer above those facts: an **interpreted, action-oriented reading of the customer’s decision** that combines engagement, consumption of the explanation, questions asked, and response actions into a single assessment of where the customer is and what the business should do next. Each major category was built around a different object — and none of those objects is the customer’s decision process as an interpreted whole. That is the gap this paper is about, and it is a gap of *interpretation and recommended action*, not of raw visibility.

4.1 What each system was built to manage

System category	Post-quote events it can capture	The interpreted decision picture it does not provide
CRM (e.g. HubSpot, Salesforce)	Quote sent, viewed, downloaded, printed, signed; activity logged on a timeline.	A synthesis of those events into a reading of intent and a prioritised next action for this specific quote, today.
Field-service / job management (e.g. ServiceM8)	Online quote viewing and acceptance; the quote’s status against the job.	Whether a still-open quote is being actively weighed, has stalled, or is effectively dead — and what to do about it now.
Quoting & accounting (e.g. Xero, Zoho Books)	Accept, decline (with reason), comment or question; Draft → Sent → Accepted / Declined → Invoiced.	Interpretation of the silence between Sent and a decision — the period where most of the uncertainty actually lives.
Follow-up / sequence automation	Scheduled reminders sent on a fixed cadence.	Whether a given reminder is warranted; it fires on a timer, not conditioned on the customer’s actual engagement.

Read down the right-hand column and the gap comes into focus. It is not that these systems see nothing — the left column shows they see a good deal. It is that what they see arrives as *disconnected facts scattered across different tools*, each owned by a system built around a different object. CRMs manage the relationship. Job tools manage the job. Quoting tools generate the document and log its status. The customer’s decision process, considered as a whole and turned into a recommended action, is not any of their jobs — it falls in the seam between them. That seam is where the interpreted intelligence has to be assembled, and no incumbent category was designed to assemble it.

4.2 Why the events alone are not enough

Even the events that are captured are noisier than they look, which is precisely why raw events cannot substitute for interpretation. Two well-documented mechanisms distort the most basic signal, the email or quote “open.” Apple’s Mail Privacy Protection, by Apple’s own account, downloads remote content in the background when a message is received rather than when it is read, and hides the recipient’s IP address — so a recorded open may not represent a human action, while genuine reading behaviour may no longer be distinguishable accurately.^[3] Independently, Microsoft Defender for Office 365’s Safe Links scans URLs during mail flow and again at time of click — so automated security infrastructure may access or inspect a tracked link without the same meaning as a deliberate customer click.^[4] A single open or click, in short, is a weak and easily corrupted signal.

The deeper point is not that the events are worthless — they are useful once interpreted — but that a data point is not a decision. Knowing a quote was opened does not tell an operator whether the customer understood the scope, where the business sits against competitors, or whether today is the day to call. Turning “opened, twice, yesterday” into “this one is worth a call this afternoon” requires a layer that weighs signal strength, combines multiple behaviours, and outputs an action. That layer is what the incumbent categories were not built to provide.

There is also an architectural reason the interpretation is not simply bolted on. Producing it well requires instrumenting the *customer's* experience of the quote — the document, the explanation, the questions, the actions available to them — and synthesising what happens there into a reading fed back to the business. A system organised around the business's own internal records (its ledger, its job cards, its relationship history) is oriented toward a different object, and adding a genuine decision-interpretation layer is closer to a new capability than a new field. The gap is not an oversight in existing products. It is a consequence of what those products were designed to be.

5. The Cost of the Unmanaged Phase

Because the post-quote phase is rarely captured as an interpreted process, its cost is largely undocumented — which is part of why it persists. This section assembles what can be said with confidence, drawing on industry experience and the limited available evidence. Where credible external data exists it is cited with its provenance and limitations. Where it does not, the cost is argued from workflow logic and clearly named as such rather than dressed up as a finding. Several of the propositions in this section are, in the terms of this series, hypotheses to be tested in later volumes; they are flagged as such here and gathered in the research agenda in Section 13.

5.1 Lost revenue from mistimed or absent follow-up

No rigorous study, as far as we can establish, directly measures the revenue effect of follow-up timing in the post-quote window. The nearest strong evidence comes from an adjacent context — response timing on inbound enquiries — and it must be read as adjacent, not as proof about quotes. The most-cited work is the Lead Response Management study led by James Oldroyd (MIT Sloan, with InsideSales.com, 2007), which examined more than 15,000 web-generated leads and over 100,000 call attempts across six companies. Its durable finding, in the study's own terms: the reported *odds* of contacting a lead were roughly 100 times higher when the first attempt came at five minutes rather than thirty, and the reported odds of qualifying a lead were roughly 21 times higher over the same interval.^[5]

The pattern's direction was reproduced — though by a different method — in a 2011 *Harvard Business Review* analysis (Oldroyd, McElheran, Elkington, “The Short Life of Online Sales Leads”), which audited 2,241 US companies and found the average firm took on the order of two business days to make first contact with an inbound lead.^[6] Two caveats must travel with these figures. First, the 2007 study draws on one vendor's platform data rather than a randomised controlled trial. Second, and more important here, both studies concern the *inbound-enquiry* phase — first contact after a customer raises their hand — not follow-up after a price has been delivered. They examine neither quotes, nor service businesses specifically, nor whether behaviour-triggered follow-up improves close rates.

What the research establishes, then, is narrow but real: **in at least one adjacent sales context, customer responsiveness can decay rapidly, and timing can materially affect the likelihood of contact and qualification.** It supports — but does not prove — the proposition that timing matters in the post-quote window too, and that

the right moment is knowable only if the customer's behaviour is visible. A business that cannot see when a customer is engaging with a quote cannot time its follow-up to that engagement, and is left approximating the moment by calendar convention ("I'll chase it in a week"). Whether behaviour-informed timing outperforms that convention in the post-quote phase specifically is, at present, an open and testable question rather than a settled finding — and one this paper would rather pose honestly than overclaim.

A note on evidence. The often-quoted figures that "48% of salespeople never follow up" and "80% of sales are made on the fifth to twelfth contact," usually attributed to a "National Sales Executive Association," appear across the sales-advice web and even in print. A 2014 investigation [7] reported searching US IRS records, the Better Business Bureau, and international business registers and finding no trace of any such association — concluding that no credible source for the figures or the claimed association could be identified. We do not rely on them. Their durability, despite having no traceable origin, is itself a small piece of evidence for this paper's thesis: the post-quote and follow-up phase is so poorly instrumented that the sector has filled the vacuum with numbers no one can source.

5.2 Wasted follow-up effort

The cost is not only under-following-up. It is also *mis*-following-up: hours spent chasing quotes that were never going to convert, and reminders sent on a schedule to customers who had already decided — either way. Sequence-automation tools, by design, fire on a timer rather than on behaviour. Without visibility into whether a customer is engaged, confused, comparing, or gone, follow-up effort is allocated by guesswork. Some of it lands on live opportunities at the right moment. Much of it is spent on the wrong quotes, at the wrong time, in the wrong tone — unpaid labour with no return.

5.3 Unnecessary discounting

A common operational pattern — familiar to many operators, though not something we can put a figure to — is that silence gets read as price resistance. An operator who cannot see *why* a customer has gone quiet may assume the reason is cost, and pre-emptively offer a discount to a customer who never asked for one, and who may have gone quiet for reasons unrelated to price (a delayed decision, a confusing scope line, a partner yet to be consulted). Where this happens, margin is surrendered to fill an information gap. With visibility into the actual point of hesitation, the same operator could address the real objection without discounting reflexively. We present this as a plausible mechanism the framework is designed to reduce, not as a measured industry rate.

5.4 Poor forecasting

A pipeline of sent quotes with no interpreted signal attached is a forecast built largely on hope. If the business cannot distinguish the quote being actively weighed from the one deleted unread, every open quote tends to be weighted the same — or weighted by intuition. This flows into cash-flow planning, capacity decisions, and hiring.

The deeper problem is one of measurement integrity. A win rate — the metric most businesses reach for to judge quoting performance — is only meaningful when the business can consistently say which quotes remain genuinely active, which have quietly been withdrawn, which have expired, and which have truly been lost. When the post-quote phase is unobserved, those categories blur: a quote that died weeks ago may still sit in the pipeline as "open," and a quote under active consideration may be written off as dead. The denominator, in other words, is unreliable — and a ratio built on an unreliable denominator cannot support confident forecasting.

Better post-quote visibility improves the *integrity* of the forecast before it improves the number: it lets a business define its opportunity set honestly, so that whatever win rate it reports actually means something. We hypothesise (see Section 13) that clearer decision intelligence improves forecasting confidence; that is a proposition for later volumes to test, not a claim we can yet substantiate.

5.5 Decisions made from intuition instead of evidence

Underlying these costs is a common condition: many of the business's most valuable decisions — which opportunities to pursue, when, and how — are made largely on intuition, because interpreted evidence is not readily available. Intuition is not worthless; an experienced operator's instinct is real signal. But instinct operating without a feedback loop is hard to improve, hard to transfer to a new employee, and hard to check against what actually happened. The post-quote phase is, for many operators, the part of the sales cycle that has been left — by default rather than by decision — least instrumented.

6. The New Category: Post-Quote Intelligence

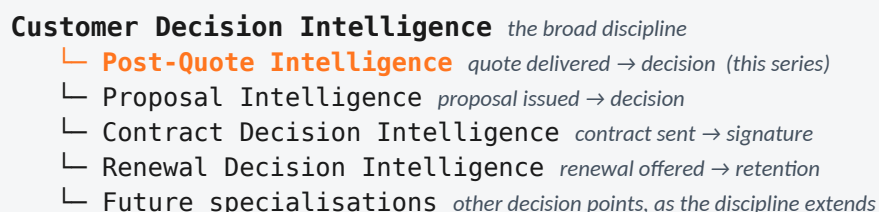
If the gap is structural, the response cannot be a feature bolted to a system built around a different object. It must be a distinct capability, defined by the phase it serves. We call that capability **Post-Quote Intelligence**, and we argue it deserves recognition as a software category in its own right — the way CRM, once a novel idea, became a category because the relationship turned out to be worth managing as a first-class object.

6.1 A definition, and a discipline

Post-Quote Intelligence is the application of Customer Decision Intelligence to the period between quote delivery and customer decision.

The broader discipline — **Customer Decision Intelligence** — is the practice of making a customer's decision process visible and interpretable to the business, so that action is taken on evidence rather than assumption. It is a general idea: any point at which a business waits on a decision it cannot see is a candidate for it. **Post-Quote Intelligence** is the first and most acute specialisation of that discipline, because the post-quote window is where the decision is highest-value, the interpreted view is most often missing, and what tooling exists tends to stop at recording events rather than reading them. Positioning it this way is deliberate: it frames Post-Quote Intelligence as the leading application of a broader family, with room for the discipline to extend to other decision points over time.

Figure 1 — Post-Quote Intelligence within Customer Decision Intelligence



Post-Quote Intelligence is the first application this series develops; the other branches are named to show the discipline's intended scope, not to claim they are addressed here.

Naming the discipline matters for a reason beyond branding. A capability without a category name is perpetually mistaken for a feature of something else — and is therefore never funded, benchmarked, or

demanding as a class. CRM had to be named before businesses could ask whether they had one. Post-Quote Intelligence is at the same stage now.

6.2 What a Post-Quote Intelligence system does

Independent of any particular product, a system in this category performs four functions that no adjacent category performs together:

1. **Instruments the customer's experience of the quote.** It makes the quote itself an observed surface — not a document that vanishes on send, but an interaction whose events can be captured.
2. **Converts behaviour into interpreted signal.** It does not merely log raw events; it translates them into a reading of where the customer is in their decision — distinguishing a strong signal from a weak or noisy one.
3. **Tells the business what to do and when.** The output is not a dashboard of vanity metrics but a next action: whether this quote warrants a call today, a wait, or a release.
4. **Writes the intelligence back into the systems of record.** It does not replace the CRM, the job tool, or the ledger. It sits alongside them and feeds them, so the decision signal lands where the business already works.

6.3 Where it fits alongside existing systems

Post-Quote Intelligence is explicitly a *layer*, not a replacement. It occupies the interpretive seam the other systems leave open — assembling their scattered events into a decision — and connects to each of them rather than competing with them:

Existing system	Relationship to the Post-Quote Intelligence layer
Field-service / job management (ServiceM8 and peers)	Owns the job lifecycle. The intelligence layer reads the quote it produced and writes decision signal back against the job — without duplicating job management.
Accounting / quoting (Zoho Books, Xero)	Owns the document and the ledger. The layer observes what the customer does with the document the ledger issued, and returns that context.
CRM (HubSpot and peers)	Owns the relationship record. The layer supplies the one thing the relationship record lacks — what the customer is doing right now with the open quote.
Follow-up automation	Owns the outbound cadence. The layer makes that cadence conditional on real behaviour, so reminders pause when the customer engages instead of firing on a timer.

This “alongside, not instead” posture is not a commercial convenience; it is a structural property of the category. The intelligence layer must connect to the systems that own the quote and the job, because those are where the decision signal has to be delivered to be acted on. A Post-Quote Intelligence product that tried to replace the CRM or the job tool would be rebuilding the wrong object and would re-create the very blind spot it exists to close.

6.4 The category in comparison

The table below sets Post-Quote Intelligence against the adjacent categories across the capabilities that matter in the decision interval. The language is deliberately cautious: individual products vary, and several incumbents offer partial versions of these capabilities. The claim is not that no other tool ever does any of this — Section 4

established the opposite — but that combining interpretation, confidence-grading and recommended action is the *defining* function of this category and only an incidental feature, at best, of the others.

Capability	Quoting software	CRM	Follow-up automation	Post-Quote Intelligence
Creates the quote	Yes	Sometimes	No	No
Records administrative status	Yes	Yes	Limited	Yes
Records engagement events	Sometimes	Sometimes	Sometimes	Yes
Interprets combined behaviour	Rarely	Limited	No	Yes
Grades confidence in intent	Rarely	Sometimes	No	Yes
Recommends the next action	Limited	Sometimes	Scheduled only	Core function
Writes insight back to systems of record	Limited	Native	Limited	Core function

Entries such as “sometimes,” “rarely” and “limited” reflect that capabilities vary by platform, plan and configuration; they are characterisations of typical category behaviour, not audits of any specific product.

7. The Five Levels of Customer Intent

A category needs a framework its users can reason with. The Five Levels of Customer Intent is a working model for interpreting what a customer's behaviour after a quote is sent may indicate — and, critically, for distinguishing strong signals from weak or ambiguous ones. It is offered as a practitioner's framework, grounded in the general principle that engagement signals form a gradient of intent rather than a binary of “opened / not opened,” and in the documented fact (Section 4.2) that the weakest signals are also the noisiest.

[3][4]

A note on evidence. The Five Levels is a reasoning framework, not an empirically validated scoring model. It orders signals by the strength of the intent they typically imply, drawing on established engagement-signal practice. It is offered to help operators interpret behaviour, not as a claim that each level converts at a fixed rate — a claim the available public data does not support and which this paper declines to invent.

The framework separates four things deliberately, because collapsing them is how the post-quote phase gets misread: the *observed event* (what actually happened), the *possible interpretation* (what it might mean), the *confidence* that interpretation warrants, and only then the *indicated action*. Each level below carries all four.

7.1 The framework

Level 0 — No observed engagement

0

No engagement with the quote has been observed. This is the most misread level: an absence of observed activity may mean the quote is genuinely unseen, or simply that engagement could not be captured. It should be treated as absence of information, not as a decision.

Observed event: No verifiable open recorded, or open-tracking unavailable or suppressed.

Confidence in the interpretation: Low. Absence of a recorded event is not evidence of disinterest.

Indicated action: Confirm receipt through a channel that does not depend on open-tracking; do not read rejection into silence.

1

Level 1 — Possible awareness

A single open with no further activity. As Section 4.2 shows, this is a weak and easily distorted signal: background image-loading and security scanners can generate opens with no human behind them, and privacy features can suppress genuine ones. It suggests possible human awareness; it does not confirm it.

Observed event: One recorded open, short or no measurable dwell, no deeper interaction.

Confidence in the interpretation: Low. Consistent with a human glance, an automated event, or noise.

Indicated action: Note it; do not over-read it. Give the customer room, and wait for a stronger signal before acting.

2

Level 2 — Apparent active review

The pattern suggests genuine engagement — returning to the quote, spending measurable time on it, moving through its detail. Repeated or sustained engagement is directionally a stronger signal than a single open, though not a certainty; it points to interest without confirming a decision is near.

Observed event: Multiple returns, sustained measurable time, movement through the document.

Confidence in the interpretation: Moderate. Repetition and duration make an automated explanation less likely.

Indicated action: A reasonable window for a relevant, well-timed contact that meets the customer while the work is apparently on their mind — offered as a judgement, not a certainty.

3

Level 3 — Deliberation or comparison

The customer engages with substance — asks a question, or (where the system can reliably observe it) attends closely to specific elements such as scope, price or terms. Interaction with content generally signals higher intent than passive review, because it takes deliberate effort.

Observed event: A question asked; or, where reliably measurable, close attention to particular sections or supporting explanation.

Confidence in the interpretation: Moderate to high for an explicit question; lower and system-dependent for inferred section-level attention.

Indicated action: Respond to the substance rather than the silence. A visible point of hesitation can be addressed directly — ideally without a reflexive discount.

4

Level 4 — Explicit decision signal

The customer takes an explicit action. This is the strongest class of signal because it is a deliberate, unambiguous act rather than an inference — though the different actions grouped here carry different predictive weight, and should not be treated as equivalent.

Observed event: An approval or accept action, a request to proceed, a callback request, or a decline. (An acceptance and a callback request are both strong, but they are not the same thing.)

Confidence in the interpretation: High — these are stated actions, not inferences — while acknowledging that a callback request is weaker evidence of a won job than an outright acceptance.

Indicated action: Act promptly and match the response to the specific action taken.

7.2 How the framework is meant to be used

The value of the model is not in the labels but in three habits it enforces. First, it forces the business to distinguish *no information* (Level 0) from *negative information* — a common and potentially costly misreading in the post-quote phase, where silence is treated as rejection when it is merely absence of visibility. Second, it grades signal strength, so effort is matched to evidence rather than spread evenly across a pipeline. Third, it turns each level into an *indicated action*, closing the loop between observation and decision that the unmanaged phase leaves permanently open.

8. A Practical Scenario

The distinction between recording events and interpreting them is easiest to see in a single worked example. The following is *illustrative* — constructed to show how the same sequence of facts reads differently through each category of tool. It is not a case study and does not report real customer data.

ILLUSTRATIVE SCENARIO — A \$48,000 COMMERCIAL HVAC QUOTE

A contractor sends a commercial HVAC quote on Monday morning.

Monday. The customer opens the quote once, briefly.

Tuesday. The customer returns twice, and this time engages with the explanation of the equipment selection.

Wednesday. The customer reviews the exclusions section closely, then requests a callback.

Consider what each category of tool shows the contractor at the end of Wednesday:

- **Quoting software shows:** “Sent.” The administrative status has not changed, because the customer has neither accepted nor declined.
- **A CRM shows:** a series of logged activities — several opens, some link clicks, a callback request — scattered on a timeline, each a discrete fact.
- **A Post-Quote Intelligence layer interprets:** active deliberation, concentrated on scope and exclusions, with an explicit next-step request — and recommends: call today, and lead with the exclusions, not the price.

The operational difference follows directly from the interpretation. Because the contractor can see that the customer’s attention settled on *exclusions* — not price — the call addresses the actual point of hesitation rather than reflexively offering a discount that was never the issue. A sales manager reviewing the pipeline can see *why* this opportunity is prioritised over others, rather than taking it on faith. And because the layer writes back to the system of record, the next person to touch the account inherits the context rather than starting cold. None of the underlying events is new; what changes is that they have been read.

9. Limitations: What Post-Quote Intelligence Cannot Reliably Know

A credible account of a category must be as clear about its limits as its claims. Post-Quote Intelligence reduces uncertainty; it does not abolish it. The following limitations are inherent, and a responsible implementation treats them as design constraints rather than inconveniences to be marketed away.

- **Not every recorded open represents a person.** As Section 4.2 sets out, automated systems can generate opens and link activity with no human behind them, and privacy features can obscure genuine ones.
- **Not all decision-making happens digitally.** A customer may decide over a phone call, a site visit, or a conversation with a partner that leaves no digital trace at all.
- **A single link may be shared among several people.** Engagement observed against one quote link may reflect two or three different readers, whose behaviours cannot be cleanly separated.
- **Dwell time can mislead.** A quote left open in a browser tab is not the same as a quote being read, and the two are hard to distinguish from timing alone.

- **Behaviour indicates probability, not certainty.** Even strong engagement is evidence of likely intent, never a guarantee of outcome.
- **Explicit actions outrank inferred intent.** An approval, a question or a callback request is firmer evidence than any pattern of passive engagement, and should be weighted accordingly.
- **Offline conversations still need to be recorded.** The intelligence is only as complete as the inputs; what happens off-platform must be captured by the operator for the picture to hold.
- **Privacy and consent requirements differ by jurisdiction.** What may be observed, retained and acted upon is constrained by law that varies by region — the subject of the next section.

These limitations do not undermine the category; they define its discipline. A system that acknowledged none of them would be making claims the evidence cannot support. The value of Post-Quote Intelligence lies precisely in interpreting uncertain signals *honestly* — grading confidence, distinguishing observation from inference, and never presenting a probability as a fact.

10. Responsible and Privacy-Conscious Use

Because Post-Quote Intelligence involves observing and interpreting customer behaviour, the category carries an obligation to define responsible use from the outset — the more so as interpretation becomes increasingly automated. The following principles are offered as a foundation for the category, not as legal advice; specific obligations under privacy and electronic-communications law vary by jurisdiction and should be assessed with qualified counsel.

- **Collect only what improves the quote and its follow-up.** Observation should be limited to what genuinely helps communicate a quote and time a follow-up — not gathered because it is technically possible.
- **Distinguish observation from inference.** What was observed (an open, a question) and what was inferred from it (likely intent) are different kinds of claim and should never be conflated.
- **Display confidence, not false certainty.** Inferred intent should be shown with its confidence level, so the operator knows how much weight it bears.
- **Do not present inferred intent as fact.** “The customer appears to be comparing” is honest; “the customer is comparing” overstates what behaviour can prove.
- **Enable transparency with customers.** Businesses should be able to explain plainly what they observe and why, rather than relying on invisible mechanisms.
- **Retain data only for legitimate operational purposes.** Signals gathered to support a live decision should not be retained indefinitely or repurposed beyond that intent.
- **Comply with applicable privacy and electronic-communications law.** Consent, disclosure and data-handling requirements differ by region and must be met wherever the business operates.

Responsible Post-Quote Intelligence should reduce uncertainty without pretending to read the customer’s mind.

Framed this way, responsible use is not a constraint bolted onto the category — it is part of what distinguishes Post-Quote Intelligence, as a disciplined practice, from mere surveillance. A layer that graded its own confidence honestly and separated observation from inference would be more useful, not less, precisely because operators could trust what it told them.

11. The Future: Decision Intelligence Becomes Standard

This closing section is openly a thesis, not a finding — a set of predictions we hold with conviction but present as the authors' view rather than as established fact. The argument runs by analogy. Categories that come to manage a genuinely valuable object have tended not to stay optional: they begin as a novel idea a few operators adopt, become an advantage for those who have them, and settle into infrastructure. Customer Relationship Management is the familiar example — within living memory, “do you have a CRM?” moved from an exotic question to an assumed one. We take this as an instructive parallel, not as proof; our view is that there is a plausible, arguably strong case that Customer Decision Intelligence — and Post-Quote Intelligence as its leading application — is early on a similar curve.

11.1 Why the direction is toward visibility

Three forces point the same way, and none of them is a passing trend:

- **Most adjacent phases now have mature tools.** Lead acquisition, quote production, job execution, invoicing, and payment have all developed mature tools for measurement and management over the past two decades. The post-quote decision phase stands out as comparatively under-instrumented. The historical direction of the sector has been toward closing gaps of exactly this kind.
- **The expectation of evidence-based decisions is increasingly widespread.** Operators who manage their marketing spend, cash flow, and scheduling on data increasingly find it anomalous to manage their most valuable decisions — which opportunities to pursue — on feel alone. Once the gap is named, running it blind reads more like a choice than a default.
- **The underlying behaviour already exists; it is only partially captured.** Customers already interact with quotes in observable ways, and some of that is now recorded — but it is typically captured in fragments and rarely interpreted as a coherent decision process. Turning those fragments into an interpreted picture is a solvable engineering and design problem, not a speculative one.

11.2 What standardisation will look like

As the category matures, the expectation will invert. Today, seeing what a customer does with a quote is a differentiator held by the few businesses that have adopted a Post-Quote Intelligence layer. Tomorrow, *not* seeing it will read the way not having a CRM reads now — as an operating business flying without an instrument its peers all use. The question a service business asks itself will shift from “why would I need to see this?” to “how did we ever decide without it?”

Our expectation is that customer decision intelligence comes to be treated as standard, much as CRM is today — not because it is fashionable, but because the decision it illuminates is too valuable to keep running in the dark.

The businesses that adopt early will not merely be more comfortable. They will make better-timed follow-up decisions, discount less reflexively, forecast more honestly, and transfer decision-making skill from the founder's instinct into a system the whole team can use. Those advantages compound. The unmanaged phase, once seen, does not go back to being invisible.

APPLIED SECTION

12. How Zevik Applies the Post-Quote Intelligence Model

The preceding seven sections describe a category, not a product. This closing section shows how one product — **Zevik** — implements the model, and is included so the framework is not left purely abstract. The category stands independently of Zevik; Zevik is presented here as a worked example of a Post-Quote Intelligence system, not as the definition of one.

8.1 What Zevik is

Zevik is post-quote follow-up software for service contractors. After a contractor sends a quote, Zevik gives the customer a clearer way to understand and respond to it, and records observable engagement associated with the recipient's quote link — including quote access, explanation playback, questions, and explicit response actions. It uses those observable signals to indicate likely intent, treating explicit actions (an approval, a question, a callback request) as stronger evidence of readiness than passive engagement, and treating a lack of activity as an absence of signal rather than a decision. It is built to run in the decision interval this paper defines, and to sit alongside the systems a contractor already uses rather than replace them.

A note on evidence. This section describes Zevik as a worked example of the Post-Quote Intelligence model, not as the definition of the category. The capabilities described reflect Zevik's approach to the decision interval; integration coverage is stated conservatively, and reflects platforms Zevik connects to rather than a claim of universal coverage.

8.2 Zevik mapped to the four functions of the category

Category function (§6.2)	How Zevik implements it
Instruments the customer's experience of the quote	Zevik turns a sent quote into an interactive experience — a clearer document, a voice explanation, questions the customer can ask, and simple response actions — making the customer's engagement observable instead of invisible.
Converts behaviour into interpreted signal	Rather than presenting raw events, Zevik indicates what the observed behaviour suggests for the contractor's next move — weighing a sustained pattern of engagement more heavily than a single open, and an explicit question or response action more heavily than passive activity.
Tells the business what to do and when	Zevik's output is a reading of which quotes appear worth pursuing and which do not, so the contractor can direct time toward apparently live opportunities rather than chasing on a fixed calendar guess.
Writes intelligence back into systems of record	Zevik connects to the contractor's existing stack — with integration centred on platforms such as ServiceM8 and Zoho Books — so the decision signal can be surfaced alongside the job and the quote rather than in a separate system the contractor has to learn.

8.3 Why the “layer, not replacement” posture matters here

Zevik does not generate quotes, manage jobs, or keep the ledger — those objects are already well served by the tools contractors rely on. It occupies the seam between them, which is precisely where this paper locates the gap. That is why Zevik is built to connect to a contractor's existing platforms rather than compete with them: the decision signal is only useful if it arrives where the contractor already works. In the terms of this paper, Zevik

is an instance of the Post-Quote Intelligence layer described in Section 6 — designed around the customer's decision, delivered alongside the systems of record.

The strategic claim of this paper is not “adopt Zevik.” It is that **existing software may record parts of the post-quote process, but service businesses that quote regularly often lack a unified capability that interprets those signals and turns them into prioritised action** — that this interpretation gap is structural rather than a matter of discipline, and that a discipline of its own, Post-Quote Intelligence, is emerging to close it. Zevik is one company's answer to that gap. The gap, and the opportunity it represents, would exist whether or not Zevik did.

13. Research Agenda: What Later Volumes Will Test

Because this is Volume I of the Post-Quote Intelligence Research Series rather than a finished account, it closes by making its own claims falsifiable. The propositions below are the working hypotheses on which the argument of this paper rests. They are stated as hypotheses precisely so that later volumes — through surveys, interviews and, in time, anonymised behavioural data — can test, refine, or refute them. Where the body of this paper has reasoned toward these claims, it has done so as argument; here they are named as what they are.

1. **Prioritisation gap.** Service businesses commonly lack a consistent method for prioritising which outstanding quotes to pursue and when.
2. **Events over interpretation.** Existing systems capture post-quote events more often than they interpret them into an assessment of intent and a recommended action.
3. **Explicit over passive.** Explicit customer actions (approval, question, callback request) are more useful predictors of outcome than passive engagement signals (opens, dwell).
4. **Behaviour over schedule.** Behaviour-informed follow-up outperforms fixed-schedule follow-up on outcomes that matter — conversion, wasted effort, or both.
5. **Visibility and forecasting.** Better post-quote visibility improves forecasting confidence and reduces wasted follow-up effort.

STATUS OF THESE PROPOSITIONS

These are research propositions, not established findings. Volume I advances them through reasoning and review of the available evidence. Later volumes will test them through original field research — and this series will report the results whether they confirm the framework, qualify it, or overturn parts of it.

13.1 How later volumes will proceed

Volume II is planned as a structured survey and interview study of service businesses, examining how quotes are actually followed up, which post-quote events are visible to operators today, how often decline reasons are captured, how live quotes are prioritised, and where existing systems fall short. Volume III, contingent on sufficient anonymised usage data, would move from stated practice to observed behaviour — examining which signals genuinely correlate with acceptance, whether repeated engagement predicts action, and whether the Five Levels framework proposed here survives contact with data or needs revision. The intent of the series is cumulative: each volume should strengthen, qualify, or correct the last, so that Post-Quote Intelligence rests over time on evidence rather than assertion.

Figure 2 — The Post-Quote Intelligence Research Series roadmap

Volume	Focus	Method
Volume I	Category definition, evidence review, research agenda	Definitional analysis and review of available evidence (this document)
Volume II	Service-business practice	Structured survey and interviews — how quotes are followed up today
Volume III	Behavioural validation	Analysis of anonymised behavioural data — which signals actually predict decisions
Annual Benchmark	Industry state-of-practice	Recurring benchmark as the evidence base matures

Volume	Focus	Method
Future reports	Thematic studies	Focused studies on specific signals, sectors or the wider discipline

Volumes II onward are planned, not yet published; their scope may be refined as the series progresses.

14. Methodology

This section documents how the sources in Volume I were selected and verified, so that the basis of the paper's claims is transparent and reproducible.

14.1 How sources were selected and verified

- **Source hierarchy.** Load-bearing claims rest on primary or official sources — vendor product and privacy documentation for current software capabilities, and original studies for research findings. Secondary and commercial sources are used only for industry interpretation or to locate primary material, and are labelled as such.
- **Product-capability claims.** Every claim about what a named platform (HubSpot, Xero, ServiceM8) does after a quote is sent was checked against that vendor's current public documentation, cited in the references. Because product capabilities change, these are stated as current at the date below and flagged for re-verification.
- **Treatment of unreliable data.** Widely circulated but unsourced figures were actively excluded; one prominent example is documented in the references as a caution rather than relied upon.
- **Where evidence is absent.** Where no reliable data exists, the paper reasons from workflow analysis and names the reasoning as such, rather than presenting inference as finding. Such points are gathered as testable hypotheses in Section 13.

14.2 Currency and review

The product capabilities and evidence described in this paper are current as of the review date on the cover (21 July 2026). Because both software features and the surrounding evidence base change, this volume is maintained rather than fixed: it will be updated when platforms add materially new post-quote capabilities, when Zevik's own research produces original data, when the Five Levels framework is tested, or when relevant privacy technology changes. A dated review cycle accompanies the series so that readers can judge how current any given volume is.

15. References and Notes on Sources

This paper follows a deliberate evidence discipline: cite real, traceable sources; state their provenance and limitations; and where reliable data does not exist, argue from workflow logic rather than supply invented figures. The numbered references below correspond to citations in the text. Several are secondary syntheses; where a finding traces to a primary study, the primary study is named.

A note on citation grade. Primary and official sources are used for load-bearing claims: Apple and Microsoft product/privacy documentation for how engagement signals are generated and distorted; the HubSpot and Xero product documentation for current post-quote capabilities; and the original Oldroyd / *Harvard Business Review* studies for response-timing evidence. Commercial and secondary sources are used only for industry interpretation or to locate primary material, and are labelled as such. Access dates reflect retrieval in 2026; readers should re-verify product capabilities, which change frequently.

- [1] HubSpot, “Manage quotes,” HubSpot Knowledge Base (accessed 2026), “Buyer activity” section. Documents the buyer actions logged as activity on a quote: quote sent, viewed, downloaded, printed, accepted, manually signed, and buyer signed (e-signed by the buyer). Specific availability varies by plan and quote type. Used to establish that CRMs now capture discrete post-quote events. knowledge.hubspot.com/quotes/manage-quotes
- [2] Xero, “Email a quote” and “Quoting software — online quotes,” Xero Central / xero.com (accessed 2026). Documents that a customer can accept, decline, and leave comments or questions through the online quote, with statuses Draft → Sent → Accepted / Declined → Invoiced and sender notification. Used to establish current quoting-tool post-quote capability. central.xero.com/s/article/Email-or-print-a-quote
- [3] Apple, “Use Mail Privacy Protection” (Apple Support) and “Mail Privacy Protection & Privacy” (Apple Legal), apple.com (accessed 2026). Apple’s own documentation states that remote content is downloaded in the background on receipt (not on open) and that the recipient’s IP address is hidden — so recorded “opens” may not reflect human action. Primary source for the noisiness of open-tracking. apple.com/legal/privacy/data/en/mail-privacy-protection
- [4] Microsoft, “Complete Safe Links overview for Microsoft Defender for Office 365,” Microsoft Learn (accessed 2026). Documents that Safe Links scans URLs during mail flow and again at time of click, meaning automated security systems can generate link-activity events. Primary source for automated distortion of click signals. learn.microsoft.com/en-us/defender-office-365/safe-links-about
- [5] Oldroyd, J. B., “Lead Response Management Study” (MIT / InsideSales.com, 2007). Analysed 3 years of behavioural data across six companies, 15,000+ web-generated leads and 100,000+ call attempts. Reported that the odds of contacting a lead were ~100× higher, and of qualifying ~21× higher, at a 5-minute vs. 30-minute first call attempt. Provenance: single-vendor platform data, not a randomised controlled trial; concerns inbound-lead response, not post-quote follow-up. Archived report PDF: [MIT Study \(archived PDF\)](#)
- [6] Oldroyd, J. B., McElheran, K., & Elkington, D., “The Short Life of Online Sales Leads,” *Harvard Business Review* (March 2011). Audit of 2,241 US companies; average time to first contact with an inbound web lead on the order of two business days. Independent-method reproduction of the 2007 pattern’s direction; likewise concerns inbound enquiries rather than quotes. hbr.org/2011/03/the-short-life-of-online-sales-leads
- [7] Rogers, S., “Those incredible sales stats everyone cites are actually completely false,” VentureBeat (15 August 2014). Reports that the “National Sales Executive Association” follow-up statistics could not be traced in IRS, Better Business Bureau or international registers, and that no credible source for the figures or the claimed association could be identified. Cited as evidence of the data vacuum in this phase; the figures themselves are not relied upon. venturebeat.com/2014/08/15/these-incredible-sales-stats-...

A closing note on method. *The reader will notice this paper contains fewer hard statistics about the post-quote phase than a document of its length might be expected to carry. That is not simply an omission; it is part of the finding — and part of the*

reason this is Volume I of a research series — The Post-Quote Intelligence Research Series — rather than a closed argument. The scarcity of rigorous public data is consistent with the proposition that the phase remains lightly instrumented and under-studied — though it does not by itself prove how any individual business manages it, since the gap may also reflect proprietary data, fragmented terminology, or research conducted under adjacent labels. Where the evidence does not yet exist, this volume reasons carefully and names its reasoning as such; later volumes will go and gather it.

16. Glossary

Terms are defined as they are used in this series, to keep the category's vocabulary consistent for readers and for the automated systems that summarise it.

Post-quote phase. The period between the delivery of a priced proposal and the customer's decision — the interval this series is about.

Decision interval. A synonym for the post-quote phase, emphasising that a decision is forming within it.

Customer intent. The customer's disposition toward accepting a quote, as estimated — never known with certainty — from observable behaviour.

Observed event. A fact that a system has recorded, such as a quote being opened or a callback being requested. Distinct from inference.

Inferred signal. A reading derived from one or more observed events, carrying a stated confidence rather than certainty.

Explicit action. A deliberate customer act — accepting, declining, asking a question, requesting a callback — which is stronger evidence than passive engagement.

Quote engagement. Observable interaction with a sent quote: opening, returning, spending time, engaging with an explanation, or acting on it.

System of record. The platform that owns a given business object — the CRM for the relationship, the job tool for the job, the ledger for the money.

Behaviour-triggered follow-up. Follow-up timed to the customer's actual engagement, as opposed to fixed-schedule follow-up sent on a timer regardless of behaviour.

Post-Quote Intelligence. The systematic capture, interpretation and operational use of customer behaviour between quote delivery and decision (see Section 1).

Customer Decision Intelligence. The broader discipline of making a customer's decision process visible and interpretable so action follows evidence; Post-Quote Intelligence is its application to the quote-to-decision window.

About the author

Izhak Salomonovitch is the founder of Zevik and has more than 25 years' experience in the building and service industries, including estimating, project delivery and commercial operations. His work on Zevik grew from the recurring operational problem this series examines — losing visibility after a quote is sent.

Disclosure: *Zevik develops software in the Post-Quote Intelligence category described in this paper. This volume is written to define and examine the category on its own terms; the applied section (Section 12) is clearly marked as the one place the paper describes Zevik's specific approach.*