

Evaluating AI-Powered Revenue Intelligence

What Actually Matters and Why. A practical framework for revenue teams making a serious technology decision.

At a Glance		
#	Section	What You'll Get
01	The Problem You're Actually Solving	Data vs. visibility vs. forecasting
02	The Nine Questions That Separate Vendors	What to ask, what the answers mean
03	How to Structure a Proof of Concept	Demo, POC, and reference call checklists
04	The Architecture Question Nobody Asks	Capture-first vs. AI-on-top-of-stale-data
05	Questions to Ask Your Own Team First	Role-specific discovery questions
06	Red Flags and Green Flags	Vendor evaluation scorecard
07	Making the Decision	Four-question go/no-go framework

Before You Start

Most evaluations of AI-powered revenue tools go sideways before they start. The demo looks great. The vendor says all the right things. The pricing fits the budget. And six months after go-live, you've got cleaner dashboards sitting on top of the same broken data.

This guide is designed to help you avoid that. It's organized around the questions that actually determine whether this kind of technology investment works, not the ones that make for a good slide deck.

One thing upfront: this guide tilts toward solutions that solve the data capture problem first. That's intentional. **Any AI applied to incomplete activity data produces confident nonsense** – and it's worth being specific about what that means. If your CRM reflects only what reps remembered to log, your AI is reasoning on a curated, partial version of reality. It will produce outputs that look precise and authoritative. Those outputs will be wrong in ways that are hard to detect, because the system has no way to flag what it doesn't know. A forecasting model that can't see 40% of a rep's activity will still give you a confidence score. It just won't be a trustworthy one. The most important question you can ask any vendor is: **where does your data come from?**

Gut check:

Remove the vendor's name and re-read their pitch. If a competitor fits just as well, you haven't found a point of differentiation. You've found a category claim.

1. The Problem You're Actually Solving

Before you evaluate vendors, get clear on the problem. Most teams walk in thinking they have a visibility problem or a forecasting problem.

CRM systems were designed for record-keeping, not signal capture. They store what reps enter – which is never the full picture. Think about what incomplete CRM data actually means downstream:

- Pipeline health scores built on partial information
- Forecast calls built on what reps believe, not what happened
- Win/loss analysis based on rep recall, not deal reality
- Coaching based on gut feel, because the evidence isn't there

A note on how teams misread this problem:

Most don't identify it as a data problem at all. They describe it as an execution problem ("our reps just don't record what we tell them to") or a systems problem ("our Salesforce instance isn't optimized"). Both diagnoses lead to the wrong fix. Tighter enforcement doesn't solve it. A new CRM configuration doesn't solve it either. The only fix is removing the manual step entirely. That's a data infrastructure problem, and it requires a data infrastructure solution.

None of these problems are solved by adding AI on top. They're solved by fixing the inputs.

Problem	What It Looks Like	What Actually Fixes It
Data Capture Gap	Activity isn't logged. Contact records are incomplete. CRM reflects what reps had time to enter.	Automatic, passive capture from email, calendar, and calls. No rep action required.
Forecast Uncertainty	Commit calls rely on rep self-reporting. Numbers move at quarter end without warning.	Signal-based validation comparing activity patterns against historical outcomes.
Coaching Blind Spots	Managers can't pinpoint why deals slip. Coaching is based on conversation, not evidence.	Activity data at the deal level – who engaged, when, how often, where momentum dropped.

2. The Nine Questions That Separate Vendors

These aren't questions to ask during a demo. They're questions to get written answers on, then verify with reference customers.

QUESTION 1

How is activity data captured?

This is the most important question you will ask. Everything else depends on it.

Ask this:

"If a rep sends 20 emails and has 4 meetings in a week and logs nothing in Salesforce, what does your system capture?"

QUESTION 2

What gets filtered out, and how?

Automatic capture sounds great until you realize not all activity is sales activity. Personal emails, internal conversations, and non-deal-related calendar events shouldn't pollute your revenue data. This is also a trust question — if reps don't trust what's being captured or how it's used, adoption will suffer regardless of how good the technology is.

- How does the system distinguish sales activity from personal or internal activity?
- Can employees configure what is and isn't captured?
- How are personal email accounts handled if reps use them for work?
- What is the process for a rep to flag something that shouldn't have been captured?

QUESTION 3

How does activity match to accounts and opportunities?

Capturing activity is step one. Associating it correctly with the right account, contact, and deal is step two — and it's harder than it sounds. A meeting with someone at a prospect company could be pre-sales, post-sales, or unrelated.

- How does the system match contacts to accounts when the contact isn't in CRM?
- How are ambiguous signals resolved — by rule, by AI, or manually?
- What's the accuracy rate on contact-to-account matching, and how is it measured?

QUESTION 4

Where does the AI actually sit in the workflow?

AI-powered revenue tools is a broad category. Some vendors use AI to summarize call transcripts. Others use it to score deal health. Others use it to answer natural language questions about pipeline. These are very different things. The question isn't whether they have AI – everyone has AI.

Ask this:

“What specific decision does your AI help a sales rep, manager, or RevOps person make that they couldn't make as well without it? Show me a real example with real output.”

QUESTION 5

How does it integrate with your existing stack?

The integration question is usually treated as a technical checkbox. It shouldn't be. The quality of the integration determines whether the system actually gets used. Sales methodology adoption improves significantly when the AI is inside the tool reps already use, not one more tab to open.

- Read-only: The vendor pulls data from Salesforce. Updates have to be manually pushed back.
- Write-back: The vendor can update Salesforce records automatically based on captured activity.
- CRM-native: The product lives inside Salesforce. Reps never leave their existing workflow.

QUESTION 6

What does implementation actually look like?

The gap between 'time to value' in the pitch and time to value in practice is where deals go wrong. A vendor who can't describe what failure looks like hasn't thought hard enough about success.

- What are the first 30, 60, and 90 days post-contract?
- Who owns the implementation – the vendor, a partner, or the customer?
- What data cleanup or CRM prep is required before go-live?
- What does a failed implementation look like, and how often does it happen?

QUESTION 7

How do you measure whether it's working?

Every vendor will offer you a list of metrics. Push past the list and ask about the baseline. You can't measure improvement without knowing where you started.

QUESTION 8

Are you paying for real value – or for features you already have?

The revenue intelligence market moves fast, and capabilities that were differentiators two or three years ago are now table stakes. Before you sign, make sure you know exactly what you're paying for. Call recording is the clearest example – Google Meet, Zoom, and Microsoft Teams all record and transcribe calls natively, at no additional cost. If a vendor's pitch leans heavily on call recording as a feature, that's a flag.

- Which of your features are also available natively in Salesforce, our sales engagement platform, or our video conferencing tools?
- What do we get from you that we genuinely cannot get from our existing stack?
- If we removed the call recording component, what would we be losing that isn't available elsewhere?

QUESTION 9

What does the security and privacy posture look like?

This section will have its own depth in your formal RFP process. But in early conversations, the questions that matter most are below. A vendor that can't answer these questions quickly and clearly has a gap in their enterprise readiness, regardless of how good the product is.

- Is customer data used to train shared models, or kept isolated?
- What are the data residency options?
- How is employee privacy handled in jurisdictions with strong labor protection laws (EU, California)?
- What certifications do you hold – SOC 2 Type II, ISO 27001, GDPR compliance?

Role	Leading Indicator	Business Outcome
VP Sales / CRO	At-risk deals identified before quarter end	Fewer late-stage losses; more accurate commits
RevOps / Sales Ops	CRM data completeness rate	Less time on manual reconciliation
Sales Manager	Activity data coverage per rep	Coaching grounded in evidence, not anecdote
CFO	Forecast variance quarter over quarter	Guidance confidence; accurate planning

3. How to Structure a Proof of Concept

Know what you're running:

A demo, a POC, and a pilot are different things with different purposes. Conflating them is how buying teams end up making decisions on insufficient evidence.

Demo: What to Watch For

Ask vendors to use anonymized real-customer data, not a curated sandbox.

- How long does it take to get from a question to an answer? Seconds matter.
- What happens when the data is ambiguous or incomplete?
- Can a sales manager use this without training? What about a rep?

Your Next Steps

- Duration defined: 30 days minimum, 60 for complex environments
- Scope is narrow: one team, one use case
- Success criteria agreed in writing before the POC starts – not after
- Using real data from your environment, not vendor-provided test data
- Running with an average team, not your most enthusiastic early adopters

The most common POC mistake:

Running it with your most tech-savvy, enthusiastic team. That's not a representative test. Run it with an average team and measure whether average people get value from it.

Reference Call Questions

Ask vendors for three reference customers. Ask those customers:

- “What did you think you were buying, and what did you actually get?”
- “What was the hardest part of implementation the vendor didn't warn you about?”
- “If you had to do it again, what would you do differently?”
- “What do your reps actually use it for on a daily basis?”
- “Has it changed how you run forecast calls? How?”

The gap between the demo story and the reference story is the most useful information you'll get.

4. The Architecture Question Nobody Asks

The problem with most sales AI isn't the AI layer. It's the data layer underneath it.

Architecture	How It Works	The Problem
A: AI on Top of CRM Data	Takes what's in Salesforce and applies intelligence to it. Better forecasting models. Smarter deal scoring.	The input data was created by reps manually entering what they remembered. Garbage in, confident nonsense out.
B: Capture First, Intelligence Second	Automatically captures all revenue activity from email, calendar, and calls. Passively. No rep action required.	Harder to build. But the AI reasons on a complete picture, not a partial one. Qualitatively different output.

The question that reveals the architecture:

“If a rep sends 40 emails in a week and logs none of them in Salesforce, what does your system know about that rep’s activity?” The answer tells you which architecture you’re dealing with.

Both architectures can produce impressive demos. The difference shows up in the deals your team doesn't catch, the forecast misses that come out of nowhere, and the coaching conversations that stay vague.

Try this in your next vendor call: **“Walk me through what happens between when a rep sends an email and when that activity shows up in your system.”** That answer will tell you more than the demo will.

5. Questions to Ask Your Own Team First

Before you talk to a single vendor, talk to the people who'll actually live with whatever you buy. The questions vary by role.

For Sales Reps

- “What information do you wish you had before a deal review that you currently don't have?”
- “How much time do you spend updating Salesforce each week, and what would you do with that time instead?”
- “When a deal slips, what's usually the first sign you see, and when do you see it?”

For Sales Managers

- “When you go into a pipeline review, what do you already know vs. what are you trying to find out?”
- “What does a deal look like at the moment you realize it's in trouble vs. the moment a rep tells you?”
- “What would change about how you coach if you had complete activity data for every rep?”

For RevOps / Sales Ops

- “What's the biggest gap between what leadership asks for in reporting and what you can actually deliver accurately?”
- “How much of your week is spent reconciling data, cleaning records, or explaining why numbers don't match?”
- “If you could change one thing about your current tech stack's data quality, what would it be?”

For Finance / CFO

- “How confident are you in the revenue forecast at the 60-day mark of a quarter, and why?”
- “When the number moves in the last two weeks of a quarter, what's your process for understanding why?”
- “What would it take for you to feel comfortable raising guidance based on pipeline data?”

6. Red Flags and Green Flags

Red Flags	Green Flags
Can't explain where their data comes from in plain language	Leads with the data capture problem before talking about AI features
Demo uses vendor-supplied data, not yours or anonymized real-customer data	Can describe a failed implementation and what they learned from it
AI capabilities are described without showing specific outputs	References use the product differently from each other – sign of real adoption
Can't name three things the product doesn't do well	Pushes back on your POC scope if it's too broad or too narrow
Customer references all have unusually similar stories	Can show outputs from a real customer environment, even if anonymized
Implementation described as 'turnkey' or 'plug and play'	Has a specific answer to 'how do we measure success in the first 90 days?'
Gets defensive when you ask about data privacy and employee consent	Is honest about which capabilities overlap with tools you already have
Leans heavily on call recording as a key differentiator – a capability now available free in most conferencing tools	

7. Making the Decision

After demos, a POC, and reference calls, it comes down to four questions. Don't move forward until you have an honest answer to each.

#	Question	Watch Out For
01	Does this solve the right problem?	Not the problem in the pitch – the problem your team actually has. If your core issue is data completeness, a better dashboard doesn't fix it.
02	Will your team actually use it?	The best technology fails with 30% adoption. Be honest about whether this fits how your team works today.
03	Can you measure impact in 90 days?	If you can't define a metric that will move in 90 days, the ROI case is speculative. Not a dealbreaker – but be honest about it.
04	What does 'no decision' actually cost?	Forecast misses, missed coaching, manual RevOps work, and late-stage losses all have a cost. The status quo isn't free.

One Last Thing

Sales AI is a real category solving a real problem. But the gap between what it can do in a demo and what it does in practice is wider than most buying teams expect – not because the technology is bad, but because the data problems underneath it are deeper than they look.

The teams that get the most value are the ones who go in with honest answers to two questions: what's actually broken, and how will we know when it's fixed?

Start there. Everything else follows.

Revenue Answers Platform