

The 2026 AI Readiness Audit for revenue teams

Is your CRM ready for AI?

We analyzed six years of B2B CRM deal history and 52 public AI readiness checklists to find what actually predicts revenue, and whether anyone measures it before AI goes live.



Your CRM already holds the signals that predict revenue, but they're measured wrong, and AI is learning from the mistakes.

AI forecasting and deal-scoring tools learn from deal history: when each deal was expected to close, how often that date moved, which stages it reached and whether it won. The quality of that history sets the ceiling on every score and forecast those tools produce.

We tested six years of B2B CRM deal history against the questions those models depend on, and audited 52 public AI readiness checklists to see who else asks them. The history held more predictive signal than the team was using, and 96% of the checklists never look for it.

SIX FINDINGS

01	Commit lost more deals than it won.	52%	of deals called Commit were lost	→
02	First forecasts almost never held.	97%	of wins missed their first-forecast month	→
03	Most pipeline outlived the sales cycle.	68%	of pipeline open 1.5x past the average win	→
04	The fields that predict wins sat empty.	2.6x	win rate when an Economic Buyer is named	→
05	Losses left a model nothing to learn.	81%	of lost deals gave no usable reason	→
06	AI now edits history nobody tests.	96%	of readiness checklists never test it	→

Commit lost more deals than it won, and almost no manager weighed in.

52%

of deals called
Commit were lost.

Commit deals still won four times
as often as the average deal.

Commit is the category leadership treats as safe, the deals a CRO puts in front of the board. Those deals did carry signal, winning four times as often as the average deal, yet more than half still closed lost. An AI forecast trained on this history learns that a Commit call is worth about half a deal.

The human check that should catch the gap was missing. The manager forecast field was blank or marked "Omitted" on 98% of open deals, so neither the team nor a model can see where a manager disagreed with a rep.


EXHIBIT 01 More than half of Commit deals and four in five Upside deals closed lost

REACHED COMMIT

EACH SQUARE IS 1% OF DEALS



52% lost **48%** won




■ Lost ■ Won □ Still open

REACHED UPSIDE

EACH SQUARE IS 1% OF DEALS



79% lost **19%** won **2%** open



Deals that reached the Commit or Upside forecast category, 2022 to 2026; a deal can appear in both. Average win rate across all closed deals: 12%. Source: AeolusGTM CRM analysis.

Reps will just progress deals even though the deal doesn't need to be progressed.

Andy Smidmore, RevOps leader, on the RevOps Lab podcast, May 2026

"All it does is give false hope to leadership," Smidmore added. When a pipeline review rewards movement, reps supply it, and the CRM records that movement as history the next model treats as fact.

CHECK YOURS

Pull every deal that reached Commit in the last four quarters and the share that closed lost. If nobody can produce that number within a day, the history an AI forecast needs isn't being kept.

First forecasts almost never held, and close dates slipped further every year.

A close date is the single field every forecast leans on, and in this history it rarely survived contact with the buyer. Nearly every won deal closed in a different month from the one first forecast, and two-thirds of all deals had their date pushed at least once.


The drift got worse over time. Deals created in 2025 averaged 143% more pushes than deals created in 2020, even though newer deals have had less time to slip. A model trained on six years of this learns to treat a moving date as normal, and stops reading the close date as a signal.

97%

of won deals missed the month they were first forecast.


Two in three deals had their close date pushed at least once.

EXHIBIT 02 Close-date pushes per deal more than doubled between 2020 and 2025




Average number of times a deal's close date moved into a later month, by the year the deal was created. Same-month moves are not counted, so these are floors. Source: AeolusGTM CRM analysis.

54% of won deals and 45% of open deals were dated on day 25 or later



Close dates by day of month, as a share of each group. The dashed line is the share an even spread would give, about 8% on days 29 to 31. Lost deals cluster less, with 19% on day 25 or later.
Source: AeolusGTM CRM analysis.

THE ONE INPUT THAT HELD UP



Actual win rate of closed deals that reached each stage. Tick marks the label.

Some month-end clustering in won deals reflects when buyers sign. Open deals carry dates a rep chose, and 45% of them sat on day 25 or later, which points to dates set to fit a reporting period rather than a buyer's timeline.

Stage probabilities told a different story, landing within six points of real win rates at every stage. The same CRM held trustworthy history and unreliable history side by side, and only testing shows which is which.

CHECK YOURS

Find the won deals that recorded a first-forecast date and count how many closed in that month. The answer tells you how much weight a forecast date deserves.

68%

of open pipeline value sat in deals open 1.5x longer than the average win takes to close.

Only 8% of deals that old ever closed won.

Most pipeline had outlived the sales cycle, and deals that old won half as often.

The average won deal closed in about four months. Two-thirds of open pipeline value had been open longer than six, and history says those deals rarely come back: past 1.5 times the average win cycle, the win rate fell by half and never recovered.

The pipeline total looks healthy because nothing ever leaves it. An AI forecast counts every one of those deals, weighted by a stage that hasn't moved in months.

CHECK YOURS

Take your average won sales cycle, multiply it by 1.5, and add up the open pipeline older than that. That share is the part of your forecast built on deals your own history says will not close.


2x

the win rate for deals that closed inside 1.5x the average cycle.

63%

of open deals carried two or more of four risk flags.


EXHIBIT 04 Win rates fell by half once a deal outlived 1.5x the average win cycle




Win rate of closed deals by sales-cycle length, as a multiple of the average won cycle. Source: AeolusGTM CRM analysis.

Nine in ten open deals carried at least one risk flag, and most carried two or more


63% / Two or more flags



30% / One flag



7% / No flags



OPEN DEALS WITH EACH FLAG

Next step untouched in 90 days 67%



No logged activity in 90 days 51%




Close date pushed 3 or more times 42%



48% of open pipeline value

Close date already passed 25%



Flags overlap, so the shares do not add up to 100%.

The fields that predict wins sat mostly empty, and the quality score never checked them.

Deals with an Economic Buyer named won 2.6 times as often as deals without one. The gap held at every stage of the pipeline, so it reflects more than late-stage deals carrying more detail. Deals with two or more buyer fields filled won 2.7 times as often as deals with none.


Only one deal in seven named an Economic Buyer, and fewer than one in fifty named a Decision Maker or a Champion. The CRM's own data-quality score graded Type, Lead Source and Next Steps instead, and the deals it rated perfect won less often than the deals it rated lowest.

2.6x

the win rate when an Economic Buyer is named.

Only 15% of deals named one.


EXHIBIT 06 Naming the Economic Buyer raised the win rate at every stage of the pipeline



Win rate of closed deals with and without an Economic Buyer recorded, overall and by the furthest stage reached. Stage results use deals with stage timestamps. Source: AeolusGTM CRM analysis.

EXHIBIT 07

Stage and close date were filled on every deal; the buyer fields on almost none



Share of opportunities with each field populated. Source: AeolusGTM CRM analysis, field profile across 19 CRM objects.

100%

of deals had a stage and close date, the two fields the CRM required.

15%

named an Economic Buyer, the field most tied to winning.

Required fields always fill, though not always with the truth. "Too often, Salesforce users get around validation rules with placeholder text like 'unknown', 'TBD', or a made-up value," Lars van Bergen wrote on [Salesforce Ben in August 2026](#).

A completeness score treats every field as equal, and the win rate shows they aren't. A readiness check worth running finds which fields predict outcomes in your own history, then measures how often your team fills them.

CHECK YOURS

Compare the win rate of closed deals with an Economic Buyer recorded against those without. If the gap is large and the field is mostly empty, your best signal is the one your reps skip.

05

Four in five lost deals gave no usable reason, so a model can't learn why you lose.


81%

of lost deals gave no usable reason.
Only 4% named a competitor.

Deal-scoring models learn as much from losses as from wins. Salesforce's Einstein Opportunity Scoring won't build a model on an org's own data until it has 200 closed-lost deals from the last 24 months, and Gong's deal model reads two years of closed deals.


Here, 42% of losses were marked "No Decision," 23% "Other" and 16% were left blank. A model reading that history can learn that deals die, and almost nothing about why.

EXHIBIT 08 Only one lost deal in five recorded a reason a team could act on



Share of lost deals by recorded loss reason. Source: AeolusGTM CRM analysis.

Each model reads three to 24 months of closed deals, losses included



- 1 Einstein Opportunity Scoring
24 months 200 won + 200 lost, plus stage history
- 2 Gong deal likelihood
24 months weighs close-date changes
- 3 Einstein Forecasting
12 months needs an update in every month
- 4 HubSpot AI forecasting
3 months closed-won dates and amounts

Sources: Salesforce, Opportunity Scoring considerations · Salesforce, Forecasting considerations · Gong, deal likelihood · HubSpot, AI projections. Accessed October 21, 2026.

89%

of logged activity was automated marketing email.

27.5%

of the remaining activity was tied to a deal.

Activity, the other input deal models read, told the same story. Nearly nine in ten logged activities were automated marketing sends, and most of the rest never attached to a deal, so the record of what reps actually did is thin exactly where a model looks for it.

CHECK YOURS

Pull last year's lost deals and count the ones with a specific reason. If most say No Decision, Other or nothing at all, your losses are teaching a model nothing.

AI can now edit deal history, and 96% of readiness checklists never test it.

HubSpot's [Smart Deal Progression](#) reads call transcripts "alongside the full deal history" and suggests changes to stage, amount, next steps and close date, and [any property can be set](#) to apply them automatically. Salesforce's [pipeline agent](#) makes the edits itself once Update Fields Autonomously is switched on.

We audited 52 public AI readiness checklists from 44 publishers, including Salesforce, HubSpot and 22 implementation partners, against 17 tests. They check access, integrations, completeness and duplicates thoroughly, and almost none test the deal history these tools learn from and now rewrite.

Of 52
readiness
checklists

79%

check access
and permissions

71%

check how
complete fields are

4%

test deal history

Two checklists test history. [Clientell's Agentforce checklist](#) sets a "Stage skip rate below 15%." [Revenue.io](#) checks stage history and is the only item that holds the CRM against an outside record, comparing logged calls with the phone system.

Eight more raise a history question without testing it. One asks whether close dates are "treated as forecast inputs." Another asks whether Commit drifts "with rep optimism and manager pressure." The tools that read live CRM data all read its current state, which shows none of the history a deal model trains on.

67%

of the checklists end in a
sales pitch.

100%


of readable quizzes score answers
you report yourself.

Hygiene tests fill most of the grid. Two checklists test deal history.

52 checklists - 44 publishers													● Tests it ○ Mentions it, no test Blank = absent					■ Tests deal history				
Data hygiene													Deal history									
Publisher / checklist	Access	Sync	Fill rate	Dupes	Format	Use case	Fresh	Usage	Stages	Links	Enrich	Comp	Close date	Stage path	Odds vs wins	Deal size	Reality					
CRM platforms																						
HubSpot / Trusted data	●	●	●	●	●		●	●	●	●	●											
HubSpot / Quality tools			●	●	●						●		○									
Salesforce / Playbook	●	●				●		●	●			●										
Salesforce / 5 ways	●	●		●	●	●	●															
Salesforce / Worksheet	●	●	●	●	●	●																
Salesforce / Gen AI	●	●	●	●	●	●																
Salesforce / Pipeline guide	●	●				●			●													
Consultancies and implementation partners																						
Aptitude8	●	●	●		●		●		●	●	●											
Babelquest	●	●	●	●	●		●			●	●											
Campaign Creators	●	●						●														
Cronyx Digital	●	●	●	●	●	●			●	●	●						○					
Customertimes / Checklist	●	●	●	●	●		●			●												
Customertimes / Guide	●	●	●	●		●	●	●	●		●											
Cyntexa	●	●	●			●	●	●	●													
DI Absolut	●	●	●	●	●	●	●	●														
EPAM	●					●			●													
Girikon	●	●	●	●	●	●		●		●												
Grazitti Interactive	●	●	●				●		●	●			○			○	○					
Huble	●		●	●	●	●		●	●													
INSIDEA	●	●	●	●	●			●	●	●		●										
IntegratelQ		●	●	●	●		●			●							○					
MO Agency	●	●	●	●	●	●	●	●	●	●	●											
Nebula Consulting	●	●				●		●														
New Breed	●	●	●	●	●																	
Perficient	●		●		●	●		●														
Royal Cyber	●		●	●	●	●	●				●											
Salesforce Break	●					●	●	●	●													
Sikich	●	●	●	●	●		●	●														
Six & Flow	●	●	●	●	●	●	●	●	●		●	●										
Squad4		●	●	●	●	●		●	●	●	●											
Synechron	●	●			●	●		●	●	●												
Trifecta						●																
Vantagepoint / Assessment	●					●																
Vantagepoint / Breeze	●	●	●	●	●	●		●	●	●	●											
Data quality tools																						
Awtana				●							●											
Clientell AI	●	●	●	●	●		●		●	●			■									
Insycle	●	●	●	●	●					●	●											
Plauti																						
Traction Complete / Checks	●	●	●	●	●	●	●	●	●	●	●											
Traction Complete / Guide										●	●											
Data and enrichment																						
Dun & Bradstreet	●	●	●	●	●		●		●	●	●											
Landbase		●	●	●	●		●			●	●											
Revenue intelligence and sales engagement																						
Airspeed (Glyphic)		●	●			●	●	●	●													
AskElephant	●		●		●		●			●												
Demodesk		●	●	●		●	●	●									○					
Nektar	●	●	●	●																		
Outreach	●	●	●	●	●	●		●	●						○							
Revenue.io	●	●	●	●	●		●	●		●			■				■					
GTM AI tools																						
Copy.ai	●	●	●	●	●	●	●	●	●		●				○							
Tofu		●	●	●	●		●										○					
Governance tools																						
Arovy	●	●																				
Research and advisory																						
Gartner	●	●				●	●	●	●													
Checklists that test it	41	40	37	33	33	29	27	25	24	20	19	2	0	2	0	0	1					

AeolusGTM audit of 52 public AI readiness checklists, October 2026. Filled squares: Clientell AI (stage history) and Revenue.io (stage history, and CRM against phone records). Hollow marks raise the idea without a test and count as no. [Coded dataset \(CSV\)](#)

EXHIBIT 11 An AI edit can become the next model's evidence



Five tests that show whether your CRM is ready for AI.

If you sell AI readiness, add these. If you buy it, ask for them by name. Each one runs on history your CRM already keeps, and the last column shows how this dataset scored.

TEST	WHAT IT MEASURES	HOW THIS DATASET SCORED
Commit outcomes	Share of deals that reached Commit and then closed lost.	FAIL 52% lost
First-forecast accuracy	Share of won deals that closed in the month they were first forecast.	FAIL 3% held
Pipeline age	Share of open pipeline value older than 1.5x your average won cycle.	FAIL 68% of value
Buyer-field signal	Win rate with and without key buyer fields recorded, and how often they're filled.	SIGNAL FOUND 2.6x lift, 15% filled
Loss-reason quality	Share of lost deals with a specific, actionable reason.	FAIL 19% usable

One input passed: stage probabilities landed within six points of real win rates at every stage, which shows a CRM can hold trustworthy and unreliable history at the same time.

Three moves before an AI forecast reads your CRM.

1

Keep AI off the forecast fields.

In HubSpot, leave the Automated box unchecked for close date, stage and amount. In Salesforce, leave Update Fields Autonomously inactive. Check a month of AI suggestions against what closed before you change either.

2

Measure what Commit is worth.

Pull the share of Commit deals from the last four quarters that closed lost. In this dataset it was 52%. If it takes more than a day to produce, fix that before a model reads the same records.

3

Find the fields that predict your wins.

Compare win rates with and without Economic Buyer, Budget and Decision Maker recorded. Require the ones that move win rate at the stage where reps learn them.

THE AEOLUSGTM CRM DIAGNOSTIC

We run all five tests on your CRM and show you which records AI can trust.

It connects read-only to HubSpot and scores data quality, pipeline hygiene and the history tests in this report: Commit outcomes, first-forecast accuracy, pipeline age, buyer-field signal and loss-reason quality.

Then we walk you through the deals behind each result and show you which fields predict wins in your own history, which records an AI forecast should ignore, and what to fix first.

You connect

HubSpot, read-only. For any other CRM, we work from an export.

You receive

A scored report, a live walkthrough and a ranked list of what to change before you trust an AI forecast.

If it's worth doing

That list becomes the scope of the work. If not, the report is yours.

[Book the diagnostic](#)

30 minutes with Sam or Chase. No deck.

Checklist audit

In October 2026 we searched the web for public AI readiness checklists, assessments and guides that tell CRM and revenue teams what to check before using AI. 52 items from 44 publishers qualified. Each was coded yes or no on 12 hygiene dimensions and five history tests, listed in the CSV. A test counts only when the item contains an explicit check; eight borderline items count as no. Seven gated items were coded on their public portion.

Vendor documentation

Read October 21, 2026, from Salesforce, HubSpot and Gong help centers, implementation guides and newsrooms.

CRM data

An anonymized B2B CRM dataset covering March 2020 to April 2026: 1,601 opportunities, 733,834 records and 1,021,915 pipeline snapshots. Win rates use closed opportunities. Stage results use opportunities with stage timestamps. Commit and Upside cover 2022 onward. Sales-cycle comparisons use the average won cycle. First-forecast accuracy uses won deals with a recorded first-forecast date. Pushes count close-date moves into a later month. Open-pipeline figures are as of May 2026.

Our interest

AeolusGTM sells CRM diagnostics and revenue-system work, and the CRM data comes from that work. Results are descriptive; we ran no significance tests. The checklist dataset is public so you can check every call we made. If we miscoded yours, [message Sam on LinkedIn](#) and we'll correct it.