REVBLACK

WIN-LOSS ANALYSIS AGENT CHECKLIST

Find out why you really win and lose deals, and what to do about it, in what order.

WHO IT'S FOR

Salesforce admins and RevOps leads, plus the CRO or sales leader who wants the real reasons behind closed deals.

WHAT YOU NEED

Salesforce Enterprise Edition or higher, 30 or more closed deals from the last 12 months, and a read-only user.

FULL GUIDE

revblack.com/guides/ai-win-loss-analysis-agent

How to use this checklist

Work through the 13 steps in order. Each one starts with what you'll have when it's done, so you know what you're working toward. Tick each box as you go.

The agent reads each closed deal in Salesforce and writes the same nine-field card for every one: the outcome, the deal story, the primary reason, secondary factors, the competitor, objections, pricing discussion, champion status and a confidence level. Instead of "Timing" or "Price" on most of your lost deals, you get why you won or lost, who you lost to and what the buyer pushed back on, and every finding traces back to named deals.

1 **Phase 1, steps 1 to 7**
Run the agent across your last 12 months of closed deals on your own computer, and turn what it finds into decisions. Nothing is written to Salesforce.

2 **Phase 2, steps 8 to 13**
Build it inside Salesforce so every deal gets a card as it closes, then pilot it and start the digests.

Check these three things first

If one of these fails, stop and fix it, because everything after depends on it.

- [ ] You have at least 30 closed deals from the last 12 months. Run an Opportunity report filtered to closed deals in that window and read the count. Below 30, any pattern is noise.
- [ ] Lost deals are really marked Closed Lost. If they get deleted or left open forever, there is nothing to analyze.
- [ ] Your org is Salesforce Enterprise Edition or higher. Check Setup under Company Information. Professional lacks API access without the Web Services API add-on.

PHASE 1 · STEPS 1 TO 4

LEARN FROM YOUR LAST 12 MONTHS

1 Set up the template

When this step is done, you'll have the prompt that fills in the nine fields for every closed deal.

- [ ] Copy the card prompt on page 8 into a file. It is a list of fields you give Claude, so nothing gets built in Salesforce.
- [ ] Replace the bracketed parts: company, product, ideal customer, sales stages, known competitors and your closed-lost picklist values.
- [ ] Keep the keys exactly as written, because the pattern pass and Phase 2 both depend on them.

2 Pull the baseline numbers

When this step is done, you'll have your baseline numbers and the data quality facts about your own records.

- [ ] Ask your Salesforce admin for a read-only user, then log in with `sf org login web` as that user. If you can't get access, ask the admin to run the reports and send you the exports.
- [ ] Export every Opportunity closed in your window, with its stage history, as closed_opps.csv and opp_history.csv.
- [ ] Ask Claude for baseline_cuts.xlsx, one sheet per cut: win rate overall and by segment, industry, deal size band, lead source, type and rep; cycle length and time per stage; the stage each lost deal died in; closed-lost reasons; competitors.
- [ ] Check the numbers: the closed deal count matches a Salesforce report, the win rate matches a report grouped by Won, and one rep's numbers match what their manager expects.
- [ ] Write down the share of lost deals with no reason or "Other", with zero logged activity and with no contact role, and whether emails and transcripts were visible at all.

3 Pull the deal record and transcripts

When this step is done, you'll have a clear picture of which emails and call transcripts the agent can read.

- [ ] Export the activity layer and contact roles the same way, as tasks.csv, events.csv, emails.csv and contact_roles.csv.
- [ ] Check that emails live in Salesforce as EmailMessage records. If they don't, run without them and flag the gap in your findings.
- [ ] If you have call transcripts, save each one as a text file named after its Opportunity ID, in a folder called transcripts.
- [ ] Finish the data quality facts from step 2: how many deals have no activity, no contact role or no emails.

4 Write the business questions

When this step is done, you'll have three to five business questions your sales leader has agreed on.

- [ ] Draft three to five questions yourself. Typical ones: who are we losing to and where, does win rate differ by segment or lead source, what do won deals have that lost ones don't, is pricing the problem, which reps close what and why.
- [ ] Take them to your sales leader and ask what to add or remove, rather than asking the team to write them from scratch.

TURN IT INTO DECISIONS

5 Run the prompts

When this step is done, you'll have a card for every closed deal in cards.jsonl and a written answer to each business question.

- [ ] Run the card prompt once per deal. Up to about 150 deals, start Claude Code in the folder with your exports and have it do the first 5, show you, then continue.
- [ ] Above 150 deals, or for a rerun, ask Claude to write a script that sends one request per deal through the Message Batches API.
- [ ] Run the pattern prompt on page 9 with cards.jsonl, baseline_cuts.xlsx and your questions from step 4.
- [ ] Require named deal examples behind every claim.

6 Check the agent's work

When this step is done, you'll have spot-checked 10 cards and know whether the prompt needs another pass.

- [ ] Pick 10 cards at random and read the underlying activity yourself.
- [ ] If more than 2 are wrong or overstated, tighten the prompt and rerun the per-deal requests.
- [ ] Label every competitor as detected in your findings until someone confirms it.

7 Turn the patterns into decisions

When this step is done, you'll have your findings turned into decisions, ordered by revenue impact, with the deal count behind every claim.

- [ ] Turn each finding into one decision: change pricing or packaging for a segment, stop chasing a segment, add to the battlecard, coach reps on an objection, or rebuild the closed-lost picklist.
- [ ] Order the decisions by revenue impact, sized by the lost Amount in the affected segment. Mark anything resting on fewer than 10 deals as directional, and say how you will measure it after 90 days.
- [ ] Give the deals where the rep's closed-lost reason disagrees with the activity their own line. It is the clearest argument for running the agent on every deal.
- [ ] Present to sales leadership and write down which recommendations they accept and who owns each one.

BUILD IT IN SALESFORCE

8 Clean up the picklist and choose where it runs

> **When this step is done,** you'll have a closed-lost picklist the agent and the reps both use, and a decision on where the agent runs.

- [ ] If Phase 1 showed it needs it, clean up the closed-lost picklist: fewer, clearer values and no "Other".
- [ ] Choose where the agent runs. A Prompt Builder template is the default. An Agentforce agent is worth it if you already own the add-on. Claude through Salesforce's hosted MCP servers gives you the most control over the prompts.

9 Keep the agent inside its own fields

> **When this step is done,** you'll have an agent that writes only to its own fields, with the rep confirming every suggested closed-lost reason.

- [ ] Create six fields on the Opportunity: AI_Win_Loss_Summary__c (Long Text Area), AI_Suggested_Loss_Reason__c (Text), AI_Competitor_Detected__c (Text), AI_Key_Factors__c (Long Text Area), AI_Analysis_Date__c (DateTime) and Needs_Win_Loss_Analysis__c (Checkbox).
- [ ] Put them in a Win-Loss Analysis section on the page layout, make them read-only for reps and managers, and set field-level security so only the integration user can edit them.
- [ ] If your org has no closed-lost reason or competitor fields, create them as picklists using the values you agreed in Phase 1.

10 Build the trigger and the template

> **When this step is done,** you'll have a flow that flags each deal when it closes and a template that writes the card to the deal.

- [ ] In Setup, create a **Record-Triggered Flow** on Opportunity that runs when **A record is updated**, with Closed equals True and **Only when a record is updated to meet the condition requirements** selected. Set **Optimize the Flow for** to **Actions and Related Records**.
- [ ] Add an **Update Triggering Record** element that sets Needs Win-Loss Analysis to True. Test it by moving a test Opportunity to Closed Lost, then activate it.
- [ ] Confirm Foundations, Data 360 and Einstein generative AI are on. In Prompt Builder, create a **New Prompt Template** of type Flex, with your card prompt as the instructions, grounded on the Opportunity plus a small flow that gathers activity, emails, contact roles and transcript.
- [ ] Give the Default Workflow User, set in Process Automation Settings, the Prompt Template User permission.
- [ ] On an asynchronous path in the flow, add the Prompt Template action. After it, call a second flow, Write Win-Loss Fields, that saves the card into the six fields and clears the flag, then notify the rep and the manager.
- [ ] To run the agent on deals that closed before the flow existed, tick the flag on them from a list view.

TEST IT, THEN KEEP IT RUNNING

11 Review the first 20 deals

When this step is done, you'll have the cards for your last 20 closed deals read through with your sales leader, with the tone, length and suggested closed-lost reasons tuned.

- [ ] Generate the cards for the last 20 closed deals with your sales leader in the room.
- [ ] Read them together, tune the tone and length, and confirm each suggested closed-lost reason matches your picklist values exactly.

12 Run the pilot

When this step is done, you'll have 15 or more closed deals marked accurate, partially accurate or wrong, a clear read on whether you've reached 80%, and a first digest reviewed with leadership.

- [ ] Run on one team or one segment for two to four weeks, long enough to catch 15 or more closed deals.
- [ ] After each batch, the manager marks every card accurate, partially accurate or wrong. Hold at 80% accurate or partially accurate before rolling out further.
- [ ] Compare the agent's suggested closed-lost reason with the rep's. Wide disagreement points to the picklist, the prompt or rep training.
- [ ] Produce the first digest and review it with leadership as a dry run of the monthly cadence.

13 Start the digests

When this step is done, you'll have both digests scheduled and an owner for acting on what they show.

- [ ] Monthly digest to sales leadership: the main reasons you won and lost, where the suggested reason disagrees with the rep's, competitors (confirmed and detected listed separately) and recurring key factors.
- [ ] Quarterly digest to marketing and product: the same over three months, plus win rate by segment, cycle length for won against lost, and the first one compares your numbers at 90 days with your Phase 1 baseline.
- [ ] On the default path, a Scheduled Flow with the **Prompt Template (Batch Generation)** action produces them. Put delivery on the **Completed** path and an error email on the **Timed out** path, and test with a real scheduled run in a sandbox.
- [ ] Put both on the calendar and agree who owns turning a finding into a change.

Getting reps on board

The main rollout risk is reps reading the agent as surveillance. Every field on the card is addressed to the deal, and none of them grade the rep. The rep and the manager see the summary at the same moment. The agent never edits the deal, and the suggested closed-lost reason is a suggestion the rep confirms. Say all of that plainly when you announce it, and give your sales leader the wording before it goes out.

Two rules that keep the findings honest

- [ ] Name the deal count behind every segment claim. Segment cuts get thin quickly, and a confident deck makes a small sample look solid.
- [ ] Write the limits down: the sample size and any missing transcripts or fields.

Where the prompts are

The next two pages hold the two prompts you run in step 5: the card prompt, which fills in the nine fields for each deal, and the pattern prompt, which answers your business questions from all the cards.

Card prompt, used in steps 1 and 5

```
You are a sales analyst producing a win-loss card for one closed Salesforce Opportunity.

Use only the material provided below. Do not infer anything the material does not support.

If the material is too thin to answer a field, write "insufficient data" in that field.

CLIENT CONTEXT
Company: [client name]. Sells [product] to [ICP]. Sales stages, in order: [list].
Known competitors: [list]. Closed Lost Reason picklist values (use these exact strings): [list].

MATERIAL
<opportunity> Name, Stage, IsWon, Amount, CloseDate, CreatedDate, Owner, LeadSource, Type,
Account.Industry, stage history </opportunity>
<activities> Tasks and Events: date, type, subject, description </activities>
<emails> EmailMessage: date, from, to, subject, body </emails>
<transcripts> call transcripts, if any </transcripts>

Return one JSON object with exactly these keys and nothing else:
{
  "opportunity_id": "...",
  "outcome": "won" | "lost",
  "story": "one paragraph, max 120 words, past tense, naming the dates and people that appear in
the material",
  "primary_reason": "one sentence",
  "secondary_factors": ["..."],
  "competitor": "name exactly as it appears in the material, or none",
  "objections": ["quoted or closely paraphrased from the material"],
  "pricing_discussed": {"yes_no": "yes" | "no", "what": "..."},
  "champion_status_at_close": "engaged" | "went silent" | "left the company" | "no champion
identified" | "insufficient data",
  "suggested_loss_reason": "one of the picklist values above, or null when won",
  "evidence": ["the activity, email or transcript dates that support primary_reason"],
  "confidence": "high" | "medium" | "low",
  "confidence_note": "one sentence on why, e.g. 14 activities and 2 transcripts, or 3 tasks and
no email"
}
```

Pattern prompt, used in step 5

```
You are answering [client name]'s win-loss business questions from two inputs:
1) cards.jsonl: one win-loss card per closed Opportunity, produced by the deal card prompt
2) baseline_cuts.xlsx: win rate and cycle length by segment, stage at loss, loss reason and
competitor counts

Business questions:
Q1 [e.g. Who are we losing to, at which stage, and in which segment?]
Q2 [...]
Q3 [...]

For each question:
- Answer in two sentences.
- Then give the evidence: the numbers from baseline_cuts.xlsx, with the deal count behind every
percentage, and the patterns in the cards.
- A pattern counts only if it appears in at least 5 deals and in lost deals at a clearly higher
rate than in won deals of the same segment. Give both rates.
- Cite at least two Opportunity names with close dates for every pattern.
- Where the picklist loss reason disagrees with the card's primary_reason, say so with counts.
- Keep "competitor confirmed in the material" separate from "competitor suspected".

Close with a section "What we could not answer and why" listing the data gaps: missing
transcripts, no email, segments under 10 deals.

Do not recommend anything yet.
```

REVBLACK

STUCK ON A STEP? WE'VE DONE THIS BEFORE.

RevBlack builds win-loss agents and other AI tools inside Salesforce and HubSpot. Tell us where you got to and we'll help from there.

Say hello at revblack.com/contact-us